Essays in Microeconomic Theory

Inauguraldissertation zur Erlangung des akademischen Grades eines

Doktors der Wirtschaftswissenschaften der Universität Mannheim

Xin Gao

Frühjahrssemester 2019

Abteilungssprecher: Prof. Dr. Jochen Streb
Referent: Prof. Dr. Thomas Tröger
Korreferent: Prof. Volker Nocke, Ph.D.

Tag der mündlichen Prüfung: Mai 13, 2019

To my parents

Eidesstattliche Erklärung

Hiermit erkläre ich, dass ich die vorliegende Dissertation selbstständig angefertigt und die benutzten Hilfsmittel vollständig und deutlich angegeben habe.

Mannheim, März 11, 2019 Xin Gao

Acknowledgements

I am deeply indebted to my advisor, Thomas Tröger, who has always awarded me full freedom to pursue my research interests, given me constructive feedback, and provided me financial support for academic activities. Without his continuous support and guidance, I would not have made it to this point.

I am also obliged to Volker Nocke and Martin Peitz for their incisive comments, insightful feedback, thoughtful encouragement and irreplaceable guidance in my graduate studies. I learnt decision theory from Mira Frick, who also generously supported my stay at Yale University and my job search, for which I am incredibly grateful.

I would like to thank my co-author, Ian Ball, for all the effort he devoted to our joint works. It has always been fascinating to work with him. His sharp mind has been a constant source of new insights and inspiration.

It is highly doubtful that I am able to acknowledge everyone who attended my talks, gave me valuable comments, or simply had an open door to talk. Nonetheless, I want to point out at least some of them: Hans Peter Grüner, Takakazu Honryo, Nicolas Schutz, and Emanuele Tarantino.

Part of the work in this dissertation was conducted during my visit to the Department of Economics at Yale University, for which I am very thankful. In particular, I benefited from the discussions with Dirk Bergemann, V. Bhaskar, Laura Doval, Ryota Iijima, Larry Samuelson, Takuo Sugaya, Juuso Välimäki as well as the audiences in my talks for their insightful comments and generous help during my stay in New Haven.

Special thanks go to the administrative staffs, particularly, Sandro Holzheimer, Golareh Khalilpour, Marion Lehnert, Julia Potapov, Astrid Reich, Dagmar Röttsches, Sylvia Rosenkranz, Nadine Scherer, and Claudius Werry, who guided me through various forms and made my life much easier.

Without my friends, I would have enjoyed the past few years a lot less. While the list is too long to put down everyone here, I have to say it is so wonderful that our paths have crossed.

Last but not least, I want to thank my family for their love and always being supportive that made my academic endeavor at all possible.

Contents

Introduction

My research focuses on understanding how people make decisions and what mechanism of decision-making is best for a decision maker. Good decision-making requires good information. However, information is dispersed among different parties. In many economic situations the decision maker does not have the relevant information and therefore needs to rely on the help by an expert. In this dissertation, I investigate novel aspects of two well-known decision-making protocols: To benefit from the expert's information, the decision maker can either ask him to report what he knows (communication), or she can simply let the expert make the decision himself (delegation).

Communication

In practice, decision-making is often based on both soft and hard information. While soft information can be misreported without any cost, hard information cannot be misrepresented. A decision maker may then be expected to prefer hard over soft information. Then a natural question arises: Whether and why does a receiver still solicit soft information from a sender who is known to be biased if the receiver already has access to some hard information? This is the main question addressed in Chapter 1 "*Checking Cheap Talk*" (joint with Ian Ball).

We consider a sender-receiver game in which the sender tries to persuade a receiver to purchase a good that has multiple attributes. The sender observes all the attributes and costlessly transmits a message to the receiver, who selects some (but not all) attributes to check and then chooses whether to buy. An important feature of our model is that the sender always prefers the receiver to buy the good independently of the quality, which makes credible communication between the two parties à la Crawford and Sobel (1982) impossible. However, we show that even though the sender has state-independent preferences, the sender strictly benefits from the ability to communicate. Indeed, one interpretation of the verification here is that it serves as a substitute for the missing preference alignment. If the sender has commitment power, she can further increase her utility by committing to randomize between various messages. However, we find that if the receiver can partially observe the state, cheap talk can do as well as Bayesian persuasion in some cases. We also observe that the receiver's verification has an ambiguous effect on the sender's utility unless the sender has commitment power.

Delegation

As an alternative decision-making protocol, delegation is adopted in many situations. Quite frequently, not only is decision-making delegated, but also information acquisition. However, two concerns of the decision maker arise: first, the decision maker is afraid of the expert's bias; second, the decision maker wants to encourage effort on information acquisition. This creates a trade-off for the decision maker when deciding how much discretion to give the expert if she can only specify the admissible set from which the expert can freely choose decisions but cannot provide monetary incentives. Then it is natural to ask what is the principal's optimal delegation set in this situation. This question is addressed in Chapter 2 "*Biased and Uninformed: Delegating to Encourage Information Acquisition*" (joint with Ian Ball).

We consider a game where the principal first specifies a decision set and then the agent chooses how much effort to exert to learn the payoff-relevant state. After learning the state with some probability, the agent takes a decision from the decision set prescribed by the principal. We find that the principal's optimal delegation set features a cap to restrict the agent's bias and may have a hole around "safe" decisions in order to encourage information acquisition. Interestingly, unlike in standard delegation models, the principal's payoff is maximized when the agent's bias is nonzero. In other words, the agent's bias can help the principal, as the principal can punish the agent without harming herself if the agent does not acquire information.

Communication vs. Delegation

In many situations, the motives of the expert may not be transparent to the decision maker. Also, the interactions between the decision maker and the expert are often repeated. In a long-run relationship, the expert may have an incentive to "look good," i.e., have reputational concerns. In such situations, if the decision maker can choose the decision-making protocol, should she keep control and solicit information from the expert or delegate the decision-making to the expert? I answer this question in Chapter 3 "*Reputational Cheap Talk vs. Reputational Delegation*."

I consider a two-period repeated game. In each period, the uninformed principal first decides whether to delegate the decision-making to the informed agent who is either good (not biased) or bad (biased). If she does, the agent takes an action himself. If she does not, the agent sends a cheap talk message to the principal who then takes an action. I find that in the second period, the principal is better off by keeping control instead of delegating to the agent. The first-period game can be transformed into a costly signaling game. The agent's behavior not only affects his current utility but also signals his type which affects his future utility. It turns out that the optimal authority allocation in the first period depends on a prior cut-off. If the prior about the agent being good is above this cut-off, the principal prefers delegation over communication. Otherwise, communication dominates delegation.

Chapter 1

Checking Cheap Talk

1.1 Introduction

Less informed agents often turn to biased experts for guidance. Even when the expert's private information is important for the agent's preferences, the expert's own preferences may be independent of this private information if the bias is sufficiently strong. For example, a salesperson wants a shopper to buy her product, no matter its true quality; a prosecutor wants a judge to convict the defendant, without regard to the defendant's actual guilt or innocence; a politician wants a voter's support, regardless of whether her proposal would actually benefit the voter. This level of bias makes credible cheap talk communication impossible. Indeed, the sender will send whichever message that induces the highest probability of the receiver taking the sender-preferred action.

However, in practice, the receiver can often gather additional information after receiving the sender's message. For instance, the shopper can inspect some attributes of the product, the judge can check some evidence submitted by the prosecutor, and the voter can do some research on the proposal. Then a natural question to ask is whether and why does an agent still solicit information from an expert who is known to be biased if the agent can *partially* observe the state of the world? In other words, (why) does an agent bother to listen to the biased expert for soft information (cheap talk messages) if the agent can get access to some hard information (through verification)?

In this paper, we show that with strategic partial verification, cheap talk messages by an extremely biased sender can be credible and strictly benefit both the sender and the receiver. The vital channel of influence is that the sender's message can influence *which* information the receiver seeks to acquire. The salesperson can point the shopper to the best attributes of the good; the prosecutor can guide the judge to inspect the strongest evidence against the defendant; the politician can highlight the merits of her proposal to the public.

It is well known in economic theory that if the sender is too biased, no information can be credibly communicated. To facilitate informative communication, our novel idea is to use partial verification as a remedy to the extreme conflict of interests between the two parties. In other words, the verification serves as a substitute for the missing preference alignment so that with the help of hard information, soft information from the sender becomes credible and strictly benefits both the sender and the receiver.

We obtain this result in a simple sender-receiver game. The receiver chooses between two actions and the sender has a strict, state-independent preference for one action over the other. The state is a vector of N binary attributes (i.e., either good or bad), drawn from a symmetric common prior. The sender observes the state and then sends a message to the receiver. After seeing the message, the receiver chooses some (but not all) components of the state to costlessly verify and then decides whether to buy the product at an exogenous price or not. We interpret the restriction on verification as a time or cognitive constraint, as in Glazer and Rubinstein (2004). For instance, a shopper may not bother to become an expert on all the tech specs of a smartphone; the judge may have to complete the case in a certain amount of time; the voter may only glance at a news article before voting.

The main insights can be captured by the case where the receiver can only check one attribute, which is our baseline model in Section 1.2 and 1.3. We construct a natural family of symmetric equilibria, called *top equilibria*. Under a top-k equilibrium, the sender points to the k best components of the state, without indicating their relative positions. Ties are broken by uniform randomization. The receiver then uniformly selects one of these recommended attributes to check, and buys if and only if that attribute is good. For any $1 \leq k \leq N$, we provide the sufficient and necessary condition for the existence of the top-k equilibrium, which requires the price below an upper bound. In Section 1.4, we extend the setting to the case where the receiver can check multiple attributes and we show that the same equilibria exist. Naturally, "checking one attribute" is simply a special case of "checking multiple attributes."

We discuss other equilibria in Appendix 1.7.4.

This family of equilibria has an intuitive structure. The probability of purchase is strictly decreasing in k because for smaller k, the relative quality of the checked attributes is higher and the probability of good attributes being checked is decreasing as the sender points to more and more attributes. However, the range of prices at which the equilibria can be sustained is increasing in k. For larger k, the checked attributes are more representative of the overall quality, which makes the receiver more optimistic and therefore more willing to pay a higher price. Moreover, the expected utility of the receiver upon buying is also increasing in k. As k increases, seeing a good attribute becomes a rarer but stronger signal about the quality of the product. This features an interesting trade-off between how often a good signal is observed and how strong that signal can be.

Since the sender tries to persuade an initially uninformed receiver, we compare the sender's utility under these equilibria with her utility under alternative communication structures. Specifically, we consider two information structures for the receiver—no verification and partial verification—and three communication protocols for the sender—no communication, cheap talk, and Bayesian persuasion. This gives six pairs of an information structure and a communication protocol.

We find that, even though the sender has state-independent preferences, for a range of prices, the sender strictly benefits from the ability to communicate. To illustrate this, suppose that without any information the receiver will not buy. With state-independent preferences, the sender's only objective is to maximize the probability of buying. If the receiver can randomly pick one attribute to check without talking to the sender, he will buy *only if* he sees a good attribute, which may increase the probability of buying to some extent. However, in our top-1 equilibrium, in which the receiver checks the recommended attribute after receiving the sender's message, the probability of buying significantly increases to a higher level so that all the non-zero types can sell the good with probability one. Intuitively, under partial verification without communication, seeing a good signal is purely random. But if the sender can communicate with the receiver (although the message is cheap talk), she can guide the receiver to the good attributes. Hence, we can think of the cheap talk message as a belief-coordinating device that guides the receiver to the good signals and this is incentive compatible for the sender due to the verification by the receiver.

We also find that if the sender has commitment power, as in Kamenica and

Gentzkow (2011), then she can further increase her utility by committing to randomize between various messages that induce different buying probabilities. This essentially smooths the discreteness in the equilibria arising from the sender's incentive constraints. In light of the Bayesian persuasion literature, it is not surprising that the sender can benefit from commitment. However, we show that cheap talk can do as well as Bayesian persuasion in some cases. This happens with prices such that the receiver is exactly indifferent between checking the recommended attributes and checking the unrecommended ones. Since the commitment power substantially enlarges the sender's set of communication strategies, if the receiver strictly prefers to follow the sender's recommendation when the message is cheap talk, then with commitment the sender can gradually adjust the signal structure to increase her utility until the receiver is indifferent between obedience and disobedience. However, if the receiver is already indifferent when the sender cannot commit, then the sender can only replicate the cheap talk equilibrium when she can commit.

Finally, we observe that the receiver's ability to partially verify the state has an ambiguous effect on the sender's utility unless the sender has commitment power, in which case verification can only restrict the set of posteriors the sender can induce. If the sender cannot commit so that the message is cheap talk, depending on the prior, the sender may or may not benefit from verification. Specifically, when the receiver decides to buy with his prior, then verification hurts the sender since the receiver will buy with probability less than one, as indicated in our top equilibria. However, if the receiver does not buy without any information a priori, then the sender benefits from verification since seeing a good signal makes the receiver more optimistic about the product. On the other hand, if the sender can commit to a signal structure, verification can only hurt the sender since it imposes restrictions on the sender's strategy set, whereas without verification the sender has full control over the receiver's information structure by Bayesian persuasion.

Related Literature

Crawford and Sobel (1982) introduced cheap talk games where the sender's messages are costless, unrestricted, and unverifiable. They showed that informative equilibria exist as long as the sender is not too biased. In many settings of interest, however, the sender prefers that the receiver take a particular action, no matter the state. With this level of bias, the sender cannot use cheap talk messages to credibly communicate information about the state, and hence the unique equilibrium is "babbling." If various assumptions about the cheap talk setting are relaxed, then there can be an informative equilibrium. For instance, messages can be credible if their cost depends on the state, as in Spence's (1973) classical signaling model or, more recently, in Kartik's (2009) model of cheap talk with lying costs.

Another approach, commonly referred to as "persuasion games" beginning with Milgrom (1981) and Grossman (1981), is to restrict the sender's strategy set to allow for information to be concealed but not misreported. This can be thought of as a reduced form approach to incorporate verification or evidence.

We take a different approach. The sender's strategy set is unrestricted and there is no exogenous lying cost. Instead, the receiver, after seeing the sender's message, can verify part (but not all) of the state. Instead of exogenously restricting the sender's messages, the receiver's actions, through verification, discipline the messages that the sender chooses to send in equilibrium, much in the spirit of Crawford and Sobel (1982). Indeed, one interpretation of verification is that it adds another dimension to the receiver's action space in a way that preferences are sufficiently aligned to support an informative equilibrium.

The advantage, relative to persuasion games, is that we can study the receiver's strategic decision about which information to verify. In each of the top-k equilibria, the binding deviation often involves checking one of the attributes that is not recommended. This captures an important strategic consideration that is absent in the models of persuasion. If the relevant attributes for the receiver's decision are the highest attributes, then an equilibrium can be sustained. Depending on the parameters of the environment, in other cases, the receiver would rather make his decision on the basis of the worst attributes. This introduces a new role for cheap talk. It is not about telling the receiver whether the product is of high quality, but rather *which* attributes are of high quality, so that the receiver can make his decision on the basis of those attributes.

There have been numerous extensions of cheap talk in various dimensions. The most relevant strand is the sequence of papers on multi-dimensional cheap talk (e.g., Battaglini, 2002; Levy and Razin, 2007; Ambrus and Takahashi, 2008). It is worth stressing the comparison with Chakraborty and Harbaugh (2007, 2010). What is important here is to distinguish the information structure from the preference structure. Our paper features a multi-dimensional information structure, but the preference structure is one-dimensional. Their papers can be seen as a complementary way to generate influential equilibria when the sender has state-independent preferences.

Our baseline model is rooted in Glazer and Rubinstein (2004) who study a mechanism design problem that minimizes the probability of the decision maker taking wrong action given that the information provider is biased in favor of one alternative of the decision maker. We depart from them in the following ways: On one hand, we assume away the commitment of the receiver, i.e., we study the existence and properties of a class of equilibria of the communication game rather than an optimal mechanism design problem. Although Glazer and Rubinstein (2004) show that the resulting optimal mechanism can be supported as an equilibrium outcome, i.e., commitment is not needed for the optimality of the mechanism. However, the converse is not true. In other words, in other equilibria of the game without commitment, the optimal mechanism with commitment may not be obtained. And what we are investigating belongs to the general group of equilibria which does not necessarily correspond to the optimal mechanism considered in Glazer and Rubinstein (2004). On the other hand, they take a mechanism design approach and focus on receiver-preferred equilibrium. Consequently, their focus is on extracting information from an informed sender rather than persuading an uninformed receiver. However, we are focused on a class of equilibria, and particularly the equilibrium that is sender-optimal. And we gain novel insights on whether and how the sender benefits from cheap talk, verification and commitment.

A recent extension of Glazer and Rubinstein (2004) is Carroll and Egorov (2017) who consider a similar setting (i.e., a receiver can verify only one dimension of a sender's multi-dimensional information) but focus on the range of sender's payoff functions that can support full information extraction, which is also a receiver-optimal mechanism design problem.

Another recent paper by Lipnowski and Ravid (2017) complements our work on how the sender benefits from commitment and cheap talk when the sender's preferences are state-independent. However, in their setting the receiver has no access to hard information, i.e., there is no verification.

The remainder of this paper is organized as follows. In Section 1.2, we present the baseline model in which the receiver can only check one attribute. In Section 1.3, we analyze the equilibrium and convey the main insights of this paper in this simple setting. We extend the environment to the general case where the receiver can check any number of attributes in Section 1.4. Some remarks about our model are made in Section 1.5. We conclude in Section 1.6. The proofs are relegated to Appendices 1.7.1, 1.7.2 and 1.7.3. We discuss other extensions in Appendices 1.7.4 and 1.7.5.

1.2 Model

There are two players, a sender (she) and a receiver (he). The state has $N \geq 2$ binary attributes. Formally, the state $\theta \in \Theta := \{0,1\}^N$ is drawn from a symmetric prior $\pi \in \Delta(\Theta)$ with full support. Specifically, we assume that $\theta_1, \ldots, \theta_N$ are exchangeable. The natural interpretation is that the state captures the (binary) quality of a product along N dimensions. Alternatively, the components can be interpreted as the outcomes of binary product tests, which are independently and identically distributed conditional on the (unobserved) product quality. The i-th component θ_i equals 1 or 0 according to whether the i-th attribute is good or bad. The sender's utility is 1 if the receiver buys and zero otherwise. Let $|\theta| := \theta_1 + \cdots + \theta_N$.[1] The receiver's utility from buying when the state is θ is $v(\theta) - P$ where $v : \Theta \to \mathbb{R}$ is a symmetric function such that $v(\theta) = v(\theta')$ if $|\theta| = |\theta'|$ and $v(\theta) > v(\theta')$ if $|\theta| > |\theta'|$, and P is the price of the good, which is exogenously given.[2] If the receiver does not buy, his utility is normalized to 0.

The timing is as follows. The sender observes the state realization and sends a message to the receiver. The receiver sees the message, updates his beliefs by Bayes' rule, and then costlessly checks one attribute of the state.[3] The receiver then decides whether to buy the product.

To complete the description of the model, we define strategies for both players. A *message strategy* for the sender is a function $m \colon \Theta \to \Delta(\mathcal{M})$ that maps each state

[1]We use $|\cdot|$ to denote the sum of the components of a vector. Therefore, $|\theta_B| = \sum_{i \in B} \theta_i$ for any $B \subset \{1, \ldots, N\}$. And we will refer to $|\theta|$ as the quality of the good below.

[2]We believe the exogenous price is an innocuous assumption since our main focus is on the strategic interaction between the sender and the receiver rather than how the price arises. We can think of it as the prevailing price in a perfectly competitive market.

[3]We assume the receiver has limited time or cognitive capacity for processing information as in Glazer and Rubinstein (2004) so that the receiver can check the attributes without any cost within his checking capacity but it is prohibitively costly to check attributes beyond his checking capacity. Our results still hold if the verification cost is sufficiently small, which will be discussed in Section 1.5.2.

realization to a distribution over a finite message space $\mathcal{M}$. For any message $A \in \mathcal{M}$, let $m(A|\theta)$ denote the probability that $m(\theta)$ assigns to A. Let $[N] := \{1, \ldots, N\}$. A strategy for the receiver is a pair (c, b) specifying a *checking strategy* $c\colon \mathcal{M} \to \Delta([N])$ with c_i being the probability that the receiver checks attribute i and a *buying strategy* $b = (b^0, b^1)\colon \mathcal{M} \to [0,1]^N \times [0,1]^N$ specifying the probability of buying upon seeing a bad attribute and a good attribute. Specifically, b_i^0 (respectively, b_i^1) denotes the probability that the receiver buys after checking attribute i and seeing that it is bad (respectively, good).

The payoffs of the sender and the receiver from this combination of strategies (m, c, b) are

$$u_S(m; c, b) = \sum_{\theta, A, i} c_i(A)(\theta_i b_i^1(A) + (1-\theta_i) b_i^0(A)) m(A|\theta)\pi(\theta),$$
$$u_R(m; c, b) = \sum_{\theta, A, i} (v(\theta) - P) c_i(A)(\theta_i b_i^1(A) + (1-\theta_i) b_i^0(A)) m(A|\theta)\pi(\theta)$$

respectively, where each sum is taken over all $\theta \in \Theta, A \in \mathcal{M}$, and $i \in \{1, \ldots, N\}$.

The solution concept is perfect Bayesian equilibrium.

1.3 Equilibrium Analysis: Checking One Attribute

1.3.1 Existence

First notice that, in a game of cheap talk without verification, no communication could be sustained in equilibrium. As all the cheap talk games, the babbling equilibrium always exists in which the sender sends messages that are independent of her type and the receiver just ignores the sender's message. With the state-independent preferences of the sender, this would be the only equilibrium. The reason is that no matter how coarse the messages are, the sender would always find it advantageous to send the message that induces the highest probability of buying. Thus in equilibrium, the receiver's behavior would be independent of the messages.

However, by introducing partial state verification, there can be a lot of non-babbling equilibria. We would focus on a particular family of equilibria in which the sender is playing a specific message strategy, called the *top-k* strategy, where k is a parameter in $[N]$. Under this strategy, the sender effectively "points" to the k

highest attributes of the realized state vector, without indicating any ordering among those k attributes. If there is a tie, the sender will break it uniformly. For example, suppose that $N = 3$, $k = 2$, and $\theta = (1, 0, 0)$. Then the sender claims "Attribute 1 and 2 are my two highest attributes" with probability 1/2 and "Attribute 1 and 3 are my two highest attributes" with probability 1/2. So the message space is $\mathcal{M} = \mathcal{P}_k$, where $\mathcal{P}_k$ is the set of k-element subsets of $[N]$. The top-k message strategy, denoted $m^k : \{0,1\}^N \to \Delta(\mathcal{P}_k)$, is formally defined as follows. First let

$$T_k(\theta) = \text{argmax}_{I \in \mathcal{P}_k} |\theta_I|.$$

For any $\theta \in \{0,1\}^N$,

$$m^k(A \mid \theta) = \begin{cases} 1/|T_k(\theta)| & \text{if } A \in T_k(\theta), \\ 0 & \text{otherwise.} \end{cases}$$

It will be convenient here and below to use the notation $a \wedge b := \min\{a, b\}$.[4] Notice that $T_k(\theta) = \{I \in \mathcal{P}_k \mid |\theta_I| = k \wedge |\theta|\}$. Therefore, for any $A \in \mathcal{P}_k$, we have

$$m^k(A \mid \theta) = \begin{cases} 1/\binom{N-|\theta|}{k-|\theta|} & \text{if } |\theta_A| = k \wedge |\theta| = |\theta|, \\ 1/\binom{|\theta|}{k} & \text{if } |\theta_A| = k \wedge |\theta| = k, \\ 0 & \text{otherwise.} \end{cases}$$

where $\binom{b}{a}$ is the number of combinations without repetitions of a elements from b elements.

We would like to exclude trivial equilibria in which no trade occurs. The following definition is in order:

Definition. A perfect Bayesian equilibrium is called *top-k equilibrium* if the sender is playing the top-k strategy and the probability of trading is positive.

Given the sender's top-k strategy, it is not hard to conjecture that in a top-k equilibrium, the receiver's strategy is evenly pick one attribute from those that are recommended to check and buy if and only if the result of verification is 1. Clearly, the sender is incentive compatible to honestly communicate since she is maximizing

[4] To simplify the notations below, we will assume the wedge operator "$\wedge$" has higher precedence than addition, so, e.g., $a \wedge b + c = \min\{a, b\} + c$.

the probability of buying by playing top-k strategy given the receiver's strategy. But we need to make sure the strategy specified is a best response for the receiver. The following theorem gives the sufficient and necessary condition for the existence of a top-k equilibrium. Let $\bar{S}_a^b$ ($\underline{S}_a^b$) denote the sum of a attributes chosen uniformly from b of the top (bottom) attributes.[5]

Theorem 1.1. *Let $k \in [N]$. There exists a top-k equilibrium if and only if*

$$P \leq \mathbb{E}[v(\theta) \mid (\bar{S}_1^k, \underline{S}_1^{N-k}) = (1, 0)].$$

The theorem states that there exists a top-k equilibrium if and only if the price is sufficiently low. In that case, the checking and buying strategies specified above together with the top-k strategy constitute an equilibrium. Intuitively, top-k equilibria can be sustained if the price is low enough such that the receiver will actually check the recommended attributes. If the price is too high, the receiver would rather deviate to the unrecommended attributes. But then the sender would have pointed to the worst attributes in the first place, knowing the receiver will not check them, and the equilibrium breaks down.

To better understand this theorem, we start with top-1 equilibrium to elaborate. It can be shown that when $k = 1$, the sufficient and necessary condition above can equivalently be expressed as

$$\mathbb{E}[(v(\theta) - P)\mathbb{1}(\theta \neq \mathbf{0})] \geq \mathbb{E}[(v(\theta) - P)\theta_1].$$

For necessity, suppose there exists an equilibrium of the desired form at the given price P. That is, the sender uniformly points to her highest attribute and the receiver checks the recommended attribute and buys if and only if the result of verification is 1. Then, on the equilibrium path, the receiver will buy if and only if $\theta \neq \mathbf{0}$. Moreover, the receiver must get a weakly higher payoff from obedience than from the following deviation: check attribute 1 and buy if and only if $\theta_1 = 1$. But this requirement is precisely the inequality above since the receiver's payoff from buying is $v(\theta) - P$.

For sufficiency, suppose the price P satisfies the condition. Given the receiver's strategy, the sender clearly has no profitable deviation. It remains to check that the receiver has no profitable deviation. Specifically, we need to show two kinds of deviation are not profitable for the receiver: first, the receiver ignores the sender's message

[5] To determine these top and bottom attributes, ties are broken uniformly. To simplify the notations below, set $\bar{S}_a^b = \bar{S}_b^b$ and $\underline{S}_a^b = \underline{S}_b^b$ for $a > b$.

and checks a random attribute and then buys if and only if that attribute is good; second, the receiver checks an unrecommended attribute based on the sender's message and buys if and only if it is good. By the necessity above, it is clear that the first deviation is not profitable. The difficulty lies in the second deviation since checking different unrecommended attributes induces different beliefs of the receiver. When the number of attributes is large, the second deviation can be very overwhelming to deal with.

To solve this problem, we will prove something slightly stronger than needed: even if the receiver could check the recommended attribute i and one other attribute $j \neq i$ of his choosing, his optimal strategy would still be to buy if and only if $\theta_i = 1$. Formally, we are proving optimality in the larger class of strategies that are measurable with respect to (θ_i, θ_j) for some j. In other words, we consider a hypothetical game in which the receiver can check not only one recommended attribute but also one unrecommended attribute, which gives the receiver a more general information partition structure than any information partition structure in the original game. And we show the optimality under this finer information partition structure for the receiver.

Now consider a general top-k equilibrium. Recall that $\bar{S}_1^k$ ($\underline{S}_1^{N-k}$) records the result of verification from checking the recommended (unrecommended) attributes. Notice that when the sender uniformly points to k highest attributes, the support of $(\bar{S}_1^k, \underline{S}_1^{N-k})$ is contained in the set

$$\mathbf{S} = \{(0,0), (1,0), (1,1)\}.$$

Then there is a threshold $(a, b) \in \mathbf{S}$ such that it is a best response to buy if and only if $(\bar{S}_1^k, \underline{S}_1^{N-k}) \geq (a, b)$. If $b = 0$, then this buying strategy is measurable with respect to $\bar{S}_1^k$ and hence the receiver need only check recommended attributes. If $b = 1$, then this buying strategy is measurable with respect to $\underline{S}_1^{N-k}$ and hence the receiver need only check the unrecommended attributes. Therefore, Theorem 1.1 essentially indicates that top-k equilibria can be supported if the buying threshold is low enough so that the receiver will actually only check the recommended attributes. Otherwise, the receiver would rather deviate to the unrecommended attributes.[6]

It is worth noting that among the family of top-k equilibria, top-1 equilibrium is the sender's most preferred equilibrium since the highest probability of buying is

[6]See more discussions in Section 1.4 after we present the more general result Theorem 1.2.

achieved in this equilibrium. In effect, all the non-zero types manage to sell the good with probability one if the top-1 equilibrium is sustainable.

1.3.2 Parametrization

In order to bring about the novel insights, we parametrize the model by taking $v(\theta) = |\theta|$ and imposing a specific strictly positive prior $\pi = (\pi_0, \ldots, \pi_N)$ over $|\theta|$, i.e., $\Pr(|\theta|= i) = \pi_i$ for $i = 0, \ldots, N$. Because $\theta_1, \ldots, \theta_N$ are assumed exchangeable, the probability distribution of the vector θ is pinned down by π. Specifically, $\Pr(\theta = \hat{\theta}) = \pi_{|\hat{\theta}|}/\binom{N}{|\hat{\theta}|}$ for $\hat{\theta} \in \{0, 1\}^N$. Then Theorem 1.1 reduces to the following proposition:

Proposition 1.1. *Let $v(\theta) = |\theta|$. There exists a top-1 equilibrium if and only if*

$$P \leq \frac{N\mathbb{E}(|\theta|)-\mathbb{E}(|\theta|^2)}{N(1-\pi_0)-\mathbb{E}(|\theta|)}.$$

For any given strictly positive prior π, we can find the upper bound of the price by Proposition 1.1 to sustain a top-1 equilibrium when $v(\theta) = |\theta|$. To make things more interesting, we further assume that $\mathbb{E}(|\theta|) < P$. Therefore, the receiver will not buy in the ex ante stage with the prior belief. However, the price is not necessarily below the upper bound dictated by Proposition 1.1 when it is above the ex ante expectation of $|\theta|$. The requirement $\mathbb{E}(|\theta|) < \frac{N\mathbb{E}(|\theta|)-\mathbb{E}(|\theta|^2)}{N(1-\pi_0)-\mathbb{E}(|\theta|)}$ implies $\pi_0 > \frac{\mathrm{Var}(|\theta|)}{N\mathbb{E}(|\theta|)}$.[7]

We consider two benchmarks to which we compare top-1 equilibrium with respect to the probability of buying.

Benchmark 1.1. There is only cheap talk communication but no verification.

Benchmark 1.2. There is no communication but the receiver can randomly pick one attribute to verify.

Under Benchmark 1.1, as we discussed before, there is only the babbling equilibrium and the receiver will ignore the sender's message and take an action based on his prior belief, which implies he will not buy by the assumption $\mathbb{E}(|\theta|) < P$.

Under Benchmark 1.2, note that the receiver would be determined to purchase the good *only if* the result of verification is 1. To see why, suppose the receiver checks attribute i with some probability $q_i \in [0, 1]$ such that $\sum_{i=1}^{N} q_i = 1$. If he finds the result of verification is 0, his expectation over all the attributes is $\mathbb{E}[|\theta|| \; \theta_i = 0] <$

[7]Note that neither Theorem 1.1 nor Proposition 1.1 relies on this requirement.

P.[8] Thus the receiver will not buy. Therefore, the probability of buying is *at most* $\sum_{i=1}^{N} q_i \cdot \left(\sum_{j=1}^{N} \frac{\pi_j \binom{N-1}{j-1}}{\binom{N}{j}} \right) = \frac{\mathbb{E}(|\theta|)}{N}$.[9]

Compared to these two benchmarks, the probability of buying in top-1 equilibrium is strictly higher than either of them.[10] This comparison is summarized in Figure 1.1. It is clear that even though the sender has state-independent preferences, for a range

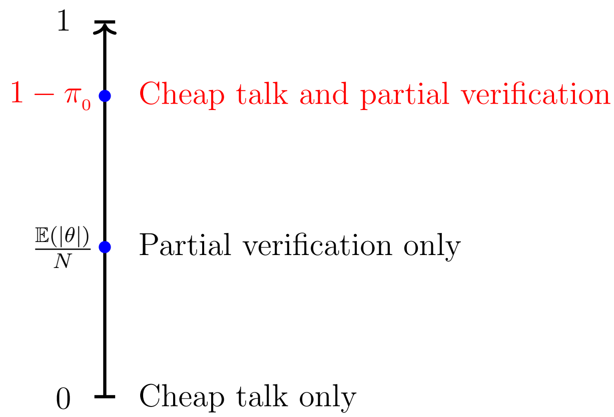

Figure 1.1. Probability of Trading

of prices, the sender strictly benefits from the ability to communicate. The underlying idea is simple. We can think of the cheap talk messages as a belief-coordinating device which guides the receiver to the good attributes. The receiver finds it optimal to follow the sender's recommendation in equilibrium and therefore he sees a good attribute more often than if he randomly picks one attribute to check and sees a "1."

1.3.3 Comparative Statics

For the following comparative statics results, we will keep N and π fixed.

[8]To see this, note that $\mathbb{E}[|\theta|| \ \theta_i = 0] \leq \mathbb{E}(|\theta|)$ (see Lemma 1.5 in Appendix 1.7.3 for the formal proof) and by the assumption that $\mathbb{E}(|\theta|) < P$, the desired inequality is attained.

[9]Since $\mathbb{E}(|\theta|) \leq \mathbb{E}[|\theta|| \ \theta_i = 1]$ (see Lemma 1.5 in Appendix 1.7.3 for the formal proof), under the assumption $\mathbb{E}(|\theta|) < P$, the price P can be higher than $\mathbb{E}[|\theta|| \ \theta_i = 1]$ so that the probability of buying can be zero even if the receiver finds the result of verification is 1.

[10]It is straightforward to see that $\mathbb{E}(|\theta|) = \sum_{i=0}^{N} i \cdot \pi_i = \sum_{i=1}^{N} i \cdot \pi_i < \sum_{i=1}^{N} N \cdot \pi_i = N(1-\pi_0)$. Therefore, we have $\frac{\mathbb{E}(|\theta|)}{N} < 1 - \pi_0$.

Equilibrium Parameter k

First, we study how the equilibrium existence depends on the equilibrium parameter k. Define the price threshold by

$$\bar{P}_k = \mathbb{E}[v(\theta) \mid (\bar{S}_1^k, \underset{\sim}{S}_1^{N-k}) = (1, 0)].$$

This is the highest price at which there exists a top-k equilibrium when the receiver can only check one attribute.

Proposition 1.2. *The price threshold $\bar{P}_k$ is strictly increasing in k.*

This means that the set of prices at which the top-k equilibrium can be sustained is increasing in k. The intuition is that as k increases, the sample of the top-k attributes is moderated so that seeing ones becomes a better signal of quality. Therefore, a positive signal from this sample can induce purchase at higher prices.

Next, we study the comparative statics of the equilibrium buying probability and the expected utility of the receiver upon buying. Let Z_k be an indicator function for whether the receiver buys. Clearly, the equilibrium buying probability is

$$\mathbb{E}[Z_k] = \tfrac{\mathbb{E}[k \wedge |\theta|]}{k}.$$

Proposition 1.3. *The equilibrium buying probability $\mathbb{E}[Z_k]$ is decreasing in k but the conditional expected utility of the receiver $\mathbb{E}[v(\theta) \mid Z_k = 1]$ is increasing in k.*

Concerning the equilibrium buying probability, as k increases, it is more and more difficult to see a good attribute for the receiver when he uniformly picks one recommended attribute to check. For example, the type (1,0,0) can sell the good with probability one in the top-1 equilibrium whereas the receiver will check her first attribute with probability 1/2 in the top-2 equilibrium. Concerning the conditional expected utility of the receiver, intuitively, if the sender points to one attribute and the receiver indeed sees a "1," he will not take it as a big deal since he suspects that the sender may only have a single "1." However, if the sender points to ten attributes and claims they are good, when the receiver sees a "1," he would reasonably believe the sender may have quite a few ones and therefore he can easily see a "1" when he randomly checks a recommended attribute. In other words, seeing a "1" becomes a rarer but stronger signal as k increases.

Proposition 1.3 implies that if the sender can select which top-k equilibrium is played, she will choose the one with the smallest k. However, she can achieve equilibria with smaller and smaller k only as the price goes down by Proposition 1.2. At some point, the price is low enough so that the receiver will buy without communication. The relative position of this threshold $\mathbb{E}[v(\theta)]$ and the thresholds $\bar{P}_k$ depends on the distribution.

Full Taxonomy

We have already seen that in terms of the probability of trading, the top-1 equilibrium is the sender's most preferred equilibrium. And the sender can strictly benefit from cheap talk even though the preferences of the two parties are misaligned. Then a natural question arises: what if a top-1 equilibrium is not sustainable because the price is too high? For example, the price is just between the price threshold of top-1 equilibrium and top-2 equilibrium so that the top-1 strategy fails to be an equilibrium. Then could the sender do better than in a top-2 equilibrium? The answer is yes. The idea is to give the sender commitment power. Here we relate to the Bayesian persuasion literature where the sender can commit to a signal structure. We will show that the sender can strictly benefit from commitment under partial verification. This may not be surprising given the insights of the Bayesian persuasion literature, e.g., Kamenica and Gentzkow (2011). However, we will also show that in some cases, the sender does not benefit from commitment and cheap talk can do as well as Bayesian persuasion.

Altogether, we consider two information structures for the receiver: no verification and partial verification (V) and three communication protocols for the sender: no communication (NC), cheap talk (CT), and Bayesian persuasion (BP). This gives six ways to pair an information structure with a communication protocol, which we label as NC, NCV, CT, CTV, BP and BPV, respectively. It is natural to compare the sender's payoff in her corresponding most-preferred equilibrium in each of these six settings. We have completed the full taxonomy. It is summarized in the following proposition.

Proposition 1.4. *The ranking of trading probability in the sender-optimal equilibrium of different settings is as follows:*

$$NC \leq CT \leq BP \text{ and } NCV \leq CTV \leq BPV \leq BP.$$

This proposition shows that no communication is weakly worse than cheap talk which is in turn weakly worse than Bayesian persuasion. This is not surprising because the sender has more and more control over the receiver's information structure and therefore she can induce higher and higher probability of buying. This is also true when we introduce verification. Since under Bayesian persuasion, the sender has full control over the receiver's information structure, she always achieves the highest probability of trading. These are the only relationships that hold in general. Any two pairs that are not related by the inequalities above have ambiguous relationships. For example, there are examples where the sender's payoff is higher under NC than NCV and also examples where the sender's payoff is higher under NCV than NC.

Numerical Example. Let $N = 3$, $v(\theta) = |\theta|$ and $|\theta|$ follows a Binomial distribution with Bernoulli parameter q where $0 < q < 1$, i.e., $|\theta|\sim$Binomial(N, q). Note that the (ex ante) expected quality $\mathbb{E}(|\theta|)$ is simply $Nq = 3q$. By Theorem 1.1, we can pin down $\bar{P}_1$ and $\bar{P}_2$ for each prior indexed by q, which are $\frac{2}{2-q}$ and $\frac{1+3q}{1+q}$, respectively. If there is no communication but the receiver can randomly pick one attribute to check, $\mathbb{E}[|\theta| \mid \theta_i = 0] = 2q$ and $\mathbb{E}[|\theta| \mid \theta_i = 1] = 2q + 1$ for $i = 1, 2, 3$ are the expected quality upon seeing "0" and "1," respectively. We plot $\bar{P}_1$, $\bar{P}_2$ and the expectations in Figure 1.2.

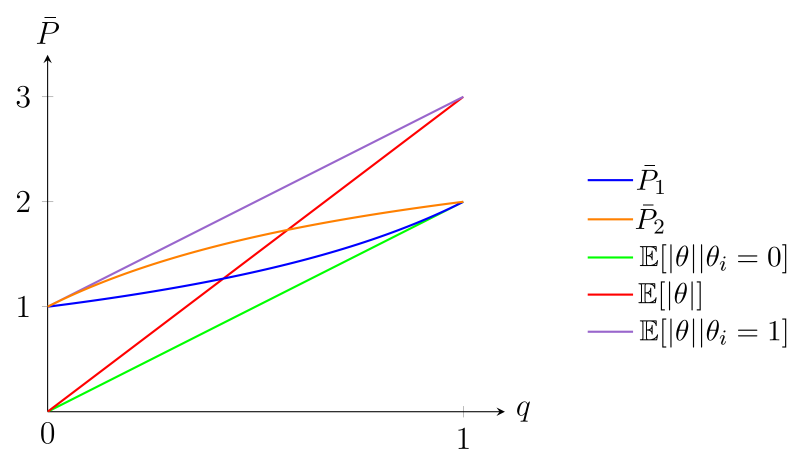

Figure 1.2. Price Cutoffs

To illustrate the potential gains from commitment power, we will restrict attention to prices $P \in (\bar{P}_1, \bar{P}_2]$. For this range of prices, we see that without communication, if the receiver randomly selects an attribute to check, he will buy if and only if it is good. For prices southeast of the red line, the receiver would buy even if he could not verify, and hence verification makes the sender worse off under no communication or cheap talk. For prices northwest of the red line, the receiver would not buy if he could not verify, and hence verification makes the sender better off under no communication or cheap talk.

Since we are more interested in the effect of the communication structure, taking verification as given, we will see an example where we have the strict inequalities NCV<CTV<BPV<BP. See Figure 1.3.

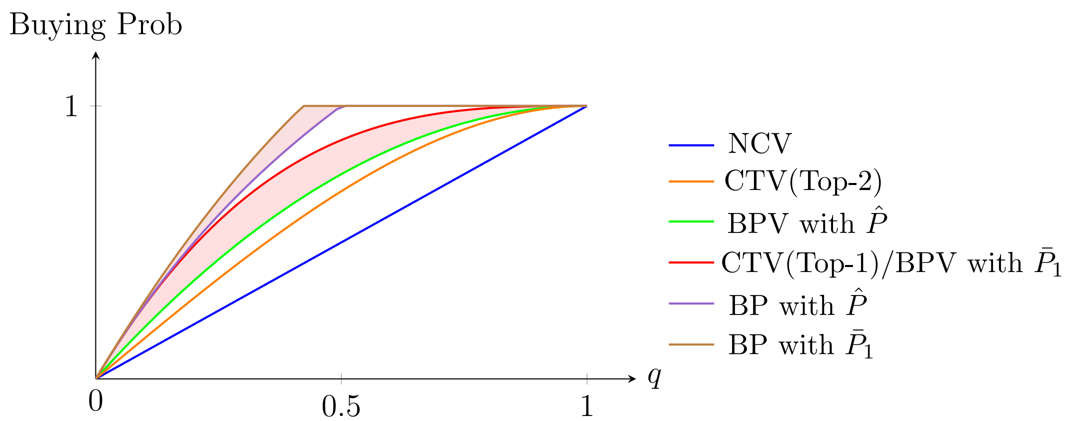

Figure 1.3. Communication Structures under Verification

We fix the price at $\hat{P}$ which is strictly between $\bar{P}_1$ and $\bar{P}_2$, specifically, $\hat{P} = 1 + q$. With this price, in the verification without communication setting (NCV), the buying probability is simply q corresponding to the blue line since the receiver will buy if and only if the result of verification is 1. In the cheap talk with verification setting (CTV), it jumps to the orange line since now only top-2 equilibrium is sustainable. When the sender has commitment power, namely, in the Bayesian persuasion with verification setting (BPV), the buying probability jumps to the green line. How? The commitment power substantially enlarges the sender's set of communication strategies. Specifically, the sender will always point to exactly one attribute, but she will mix between pointing to the highest attribute and the second highest attribute. The mixing probability will be computed by the receiver's indifference condition: the receiver will be indifferent between (i) checking the indicated attribute and buying if and only if it is good; and (ii) checking a random attribute and buying if and only if it is good. Note that at the price $\hat{P}$, the receiver strictly prefers obedience to disobedience in the top-2 equilibrium. However, under BPV, the sender can gradually adjust the signal structure to increase the buying probability and hence decrease the receiver's utility until the receiver is exactly indifferent.

Without surprise, in pure Bayesian persuasion setting, the sender's hands are not tied any more. She can induce an even higher buying probability. Specifically, the equilibrium of the BP setting will have a monotonicity property: if $|\theta|$ is larger than some marginal number, the sender will recommend the receiver to buy; if it is lower than the marginal number, the sender will recommend the receiver not to buy; and if it is equal to the marginal number, the sender will mix between the two recommendations. The marginal number and mixing probability will be pinned down by the condition that the receiver is indifferent between buying and not buying when the sender recommends she buy. It turns out that at $\hat{P}$ the marginal number is 0, which implies, now the zero type gets some chance to sell. Therefore, the buying probability is higher than any communication protocol with verification setting.

Another interesting thing is that, if we reduce the price from $\hat{P}$ to $\bar{P}_1$, the buying probability curves under BPV trace out the pink region and the buying probability curves under BP trace out the dark pink region, respectively. Under BPV, when the price varies from $\hat{P}$ to $\bar{P}_1$, the sender puts more and more weight on the highest attribute and hence the buying probability rises. When the price drops to $\bar{P}_1$, the sender cannot benefit from commitment any more under verification (CTV does as well as BPV). Under BP, when the price tends to $\bar{P}_1$, the sender will tell the receiver to buy if her type is not $\mathbf{0}$ and also more and more often when her type is $\mathbf{0}$ according to the outcome of her randomization device. And hence the buying probability rises.

1.4 Extension: Checking Multiple Attributes

Assuming that the receiver can only check one attribute provides a simple illustration of the main insights we want to bring up. However, the analysis can be generalized. In this section, we will allow the receiver to costlessly and simultaneously check any number of attributes after seeing the sender's message.[11] Of course, it is still partial verification, i.e., $n < N$ where n denotes the receiver's information (or checking) capacity.

[11] It turns out that the results are not affected if the receiver sequentially checks the attributes that he chose. See Corollary 1.4.1.

1.4.1 Existence

To define the receiver's strategy in this general setting, let $\mathcal{P}_n$ denote the family of subsets of $\{1,\ldots,N\}$ of size at most n. Here $\mathcal{P}_n$ consists of all attribute sets that the receiver can check. A strategy for the receiver is a pair (c,b) consisting of a *checking strategy* $c:\mathcal{M}\to\Delta(\mathcal{P}_n)$ and a *buying strategy*

$$b:\mathcal{M}\times\bigcup_{B\in\mathcal{P}_n}\{0,1\}^B\to[0,1].$$

A checking strategy specifies the set of attributes the receiver checks after each message. It maps each message from the sender to a distribution over subsets of attributes. Let $c(B|A)$ denote the probability that, upon receiving the message A, the receiver checks exactly the attributes in B. The checking strategy specifies the possibly random set of attributes the receiver checks upon receiving the sender's message. The realizations of these checked attributes are summarized by a function $y:B\to\{0,1\}$. And $b(A,y)$ is the probability that the receiver buys the good upon receiving message A, checking exactly the attributes $i\in B$, and observing $\theta_i=y(i)$ for each $i\in B$.[12]

Since $n=1$ is taken as a special case of what we will discuss in this section, the family of equilibria we will focus on is still the top-k equilibria. We reproduce the sender's top-k strategy and the definition of top-k equilibrium here:

Definition. The top-k message strategy, denoted $m^k:\{0,1\}^N\to\Delta(\mathcal{P}_k)$, is formally defined as follows: For any $A\in\mathcal{P}_k$,

$$m^k(A\mid\theta)=\begin{cases}1/\binom{N-|\theta|}{k-|\theta|} & \text{if } |\theta_A|=k\wedge|\theta|=|\theta|,\\ 1/\binom{|\theta|}{k} & \text{if } |\theta_A|=k\wedge|\theta|=k,\\ 0 & \text{otherwise.}\end{cases}$$

Definition. A perfect Bayesian equilibrium is called *top-k equilibrium* if the sender is playing the top-k strategy and the probability of trading is positive.

Now we define a strategy for the receiver. If he can check all the attributes recommended by the sender, he does so (and does not check any others). If he cannot check all the recommended attributes, then he checks n of them. He buys if the

[12]We will use the function notation $y(i)$ and the vector notation y_i interchangeably throughout.

number of ones he sees exceeds some threshold. We give the sufficient and necessary condition of the existence of top-k equilibrium that holds in general.

Theorem 1.2. *Suppose $n \leq N-2$. For each $k = 1, \ldots, N$, there exists a top-k equilibrium if and only if*

$$P \leq \mathbb{E}[v(\theta) \mid (\bar{S}_n^k, \underline{S}_n^{N-k}) = (n \wedge k, 0)].$$

If $n = N-1$, it goes through except the "only if" statement holds for all but finitely many prices.[13]

First note that by taking $n = 1$ Theorem 1.2 reduces to Theorem 1.1.

The theorem states that there exists a top-k equilibrium if and only if the price is sufficiently low. In that case, the checking and buying strategies specified above together with the top-k strategy constitute an equilibrium. In partibular, if the top-k equilibrium can be sustained for some $k < n$, then the receiver need not even check n of the attributes. He can base his decision on only k of them. Of course this is not a best response for the receiver in general. The intuition is that if the receiver would gain from checking unrecommended attributes, then he would check as many of the unrecommended attributes as possible. But then the sender would have an incentive to deviate and point to lower attributes, and the equilibrium would break down.

The formal proof of this theorem is relegated to the Appendix. Here we would like to sketch the basic idea of the proof. For simplicity and intuition, let $n < k$ and $n < N-k$ so that the receiver's information capacity is lower than the number of (un)recommended attributes. It seems quite overwhelming because we need to take care of quite a lot of possible deviations and different deviations induce different beliefs of the receiver. Here the idea we are using is that, since the receiver's buying strategy is based on an information partition structure induced from the checking decision and message received, we construct a hypothetical game in which the receiver can check not only n recommended attributes but also n unrecommended attributes. That is, the receiver is given a more general information partition structure and it is finer than any information partition structure in the original game. Then if we can show the receiver's buying strategy is optimal in this hypothetical game, we can conclude the receiver's buying strategy is also optimal in the original game, if it is still attainable.

[13]Specifically, when $n = N-1$ and $k \leq N-2$, we can also construct top-k equilibrium if $P = \mathbb{E}[v(\theta) \mid (\bar{S}_n^k, \underline{S}_n^{N-k}) = (n \wedge k, j)]$ for some $j = 1, \ldots, N-k-1$.

In other words, we enlarge the strategy space of the receiver, and if we can show the optimality of the receiver's buying strategy in this larger set, we can conclude it is still optimal in the smaller set if it is still attainable. Therefore, we will effectively prove something stronger than needed (we allow for more possible deviations in the hypothetical game than we need to take care of in the original game).

The key observation is that, given the sender's top-k strategy, the support of $(\bar{S}_n^k, \underline{S}_n^{N-k})$ is contained in the set

$$\mathbf{S} = \{(0,0), (1,0), \ldots, (n,0), (n,1), \ldots, (n,n)\}$$

on which the product order is total.[14] If the receiver could check n of the recommended attributes and n of the unrecommended attributes, then there is a threshold $(a,b) \in \mathbf{S}$ such that it is a best response to buy if and only if $(\bar{S}_n^k, \underline{S}_n^{N-k}) \geq (a,b)$. If $b = 0$, then this buying strategy is measurable with respect to $\bar{S}_n^k$ and hence the receiver need only check recommended attributes. Alternatively, if $b > 0$, then this buying strategy is measurable with respect to $\underline{S}_n^{N-k}$ and hence the receiver need only check the unrecommended attributes. This is the heart of Theorem 1.2.

Essentially, top-k equilibria can be sustained if the buying threshold is low enough so that the receiver will actually only check the recommended attributes. If the buying threshold is too high, the receiver would rather deviate to the unrecommended attributes. But then the sender would have pointed to the worst attributes in the first place, knowing the receiver will not check them, and the equilibrium breaks down.

It is not hard to observe that the receiver does not benefit from sequential verification.

Theorem 1.2 still holds if the receiver can check the attributes sequentially rather than simultaneously.

1.4.2 Comparative Statics

Compared to $n = 1$, not all the previous comparative statics results can be carried over to $1 < n < N$. We will discuss the reasons and give counter-examples. We still keep N and π fixed.

First, recall that the receiver's buying strategy is a cutoff strategy. That is, he will buy if the number of ones he sees exceeds some threshold. The threshold $\bar{s}^*$ for recommended attributes is defined by

[14] Recall that the product order means $(a,b) \geq (a',b')$ if and only if $a \geq a'$ and $b \geq b'$.

$$\bar{s}^*(P,k,n) = (n \wedge k + 1) \wedge \min\Big\{s \in \{0,\ldots,n\wedge k\} \mid \mathbb{E}[v(\theta) \mid \bar{S}_n^k = s] \geq P\Big\}.$$

Analogously, the threshold $\underline{s}^*$ for unrecommended attributes is defined by

$$\underline{s}^*(P,k,n) = (n \wedge (N-k) + 1) \wedge \min\Big\{s \in \{0,\ldots,n\wedge (N-k)\} \mid \mathbb{E}[v(\theta) \mid \underline{S}_n^{N-k} = s] \geq P\Big\}.$$

The following proposition gives how the buying thresholds depend on the price P, the equilibrium parameter k, and the receiver's information capacity n.

Proposition 1.5. *Both thresholds $\bar{s}^*(P,k,n)$ and $\underline{s}^*(P,k,n)$ are weakly increasing in P and n but weakly decreasing in k.*

The comparative statics in P and n are straightforward. If the price is higher, the receiver demands a more favorable signal of quality in order to buy the good. If the sample is larger, more successes are required to provide the same signal of quality. Increasing k pushes the sample of recommended attributes closer to random but also pushes the sample of unrecommended attributes away from random. In each case, the sample is made lower relative to the population as a whole, and hence fewer successes are needed to induce purchase.

Next, we study the comparative statics of equilibrium existence. Define the price threshold by

$$P^*(k,n) = \mathbb{E}[v(\theta) \mid (\bar{S}_n^k, \underline{S}_n^{N-k}) = (n \wedge k, 0)].$$

This is the highest price at which there exists a top-k equilibrium when the receiver's information capacity is n.

Proposition 1.6. *For any fixed n, the price threshold $P^*(k,n)$ is strictly increasing in k.*

The intuition is similar to the previous result. As the sender points to more and more high attributes, the receiver's sample of recommended attributes becomes more and more representative and he is more optimistic when seeing ones and therefore would like to pay a higher price.

This result is a generalization of Proposition 1.2. However, we do not necessarily have monotonicity in n. Following the equilibrium is more informative as n increases, but so is deviating. We can construct a counter-example by taking $N = 3$, $v(\theta) = |\theta|$

$P^*(k,n)$	$n=1$		$n=2$
$k=1$	1.2	$\searrow$	1
$k=2$	1.5	$\nearrow$	2

Table 1.1. Price Threshold $P^*(k,n)$

and $|\theta|$~Binomial$(3,1/3)$. By Theorem 1.2, we can calculate the corresponding price thresholds. See Table 1.1. When fixing k and varying n, there is no unambiguous monotonicity in n.

At last, with respect to equilibrium buying probability and the receiver's expected utility upon buying, we lose the preceding nice comparative statics results in Proposition 1.3 for the more general case $1 < n < N$. The fundamental reason is that the buying threshold can change when the receiver can check more than one attribute. When $n = 1$, the only equilibrium buying threshold is one.

For intuition, let's focus on the case where $n < k$ and $v(\theta) = |\theta|$. So after the sender recommends the k highest attributes, the receiver randomly and uniformly selects n of these k attributes to check. He will buy if and only if at least $\bar{s}^*$ of these n attributes are good. If $n = 1$, then we must have $\bar{s}^* = 1$. That is, the receiver will check a single attribute and buy if and only if it is good. When $n > 1$, however, we may have $\bar{s}^* > 1$. For example, if $n = 3$ and $\bar{s}^* = 2$, the receiver will check three attributes and buy if and only if at least two of them are good. Recall that $\bar{s}^*(P,k)$ is weakly decreasing in k (see Proposition 1.5).[15] If k is lower, then the receiver knows that the sample of attributes he is observing is more upwardly biased, so he demands to see more high attributes in order to be willing to buy. Also notice that no matter the value of n, if the sender points to the k highest attributes, then the average quality of these recommended attributes is weakly decreasing in k. For example, the average quality of the two highest attributes is weakly greater than the average quality of the four highest attributes.

Let's look at a simple case where $N = 4$, $n = 2$ and we can compare $k = 2$ with $k = 3$. We could get a simple failure of monotonicity as follows. For P low enough, we might have $\bar{s}^*(P,2) = \bar{s}^*(P,3) = 1$. That is, regardless of whether the sender

[15]In general, the value of $\bar{s}^*$ is also determined by the receiver's information capacity n. For fixed n, we denote it by $\bar{s}^*(P,k)$ as a shorthand.

points to the two or three highest attributes, the receiver will check $n = 2$ of them and buy if and only if at least one of those attributes is good. It follows that for this range of prices, the top-2 equilibrium induces a higher buying probability and a lower expected quality conditional on buying. But when P increases beyond some threshold, say to some P', we might have $\bar{s}^*(P', 2) = 2$ and $\bar{s}^*(P', 3) = 1$. This means that in the top-2 equilibrium, the receiver will only buy if both attributes he checks are good, but in the top-3 equilibrium, the receiver will buy if at least one of the two attributes he checks is good. In this region, where $\bar{s}^*(P', 3) < \bar{s}^*(P', 2)$, the buying probability may be higher and expected quality conditional on buying lower in the top-3 equilibrium than in the top-2 equilibrium. This phenomenon cannot happen when $n = 1$ because in that case either $\bar{s}^* = n = 1$ or else the receiver always buys.

More concretely, consider the following numerical example: Let $N = 4$, $n = 2$, $v(\theta) = |\theta|$, $|\theta| \sim \text{Binomial}(N, q)$ where $q = 0.3$ and $P \in [0.6, 1.3]$. We plot the change of the equilibrium buying probability $\mathbb{E}[Z_k]$ and the expected quality upon buying $\mathbb{E}[|\theta||\ Z_k = 1]$ as the price P varies between 0.6 and 1.3 in Figure 1.4a and Figure 1.4b, respectively. Specifically, when $0.6 < P < 1$ the buying probability is higher and expected quality conditional on buying lower in the top-2 equilibrium than in the top-3 equilibrium but the relationships are reversed when $1 < P < 1.3$.

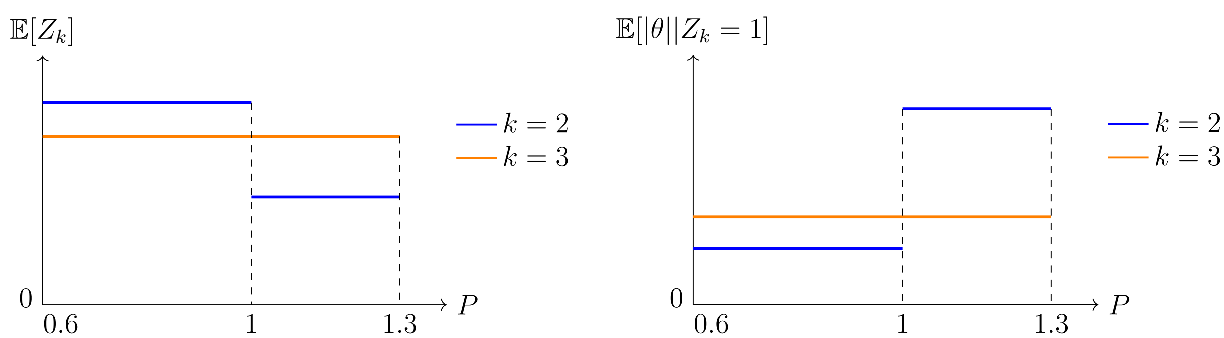

(a) Equilibrium Buying Probability

(b) Expected Quality upon Buying

Figure 1.4. Equilibrium Buying Probability and Conditional Expected Quality

1.5 Discussion

1.5.1 Sender- and Receiver-Optimal Equilibrium

Even though the state space Θ is finite and the payoff-relevant component of the receiver's action is binary, allowing for (payoff-irrelevant) partial verification means that the analysis must be performed on functions from Θ to $\mathcal{A} := \Delta([N]) \times [0,1]^N \times [0,1]^N$. Such functions form quite a large space, and this fact, together with the structure of the incentive constraints, makes it so difficult to characterize all equilibria.

In the core of the paper, we have identified a natural class of equilibria, the *top equilibria*, which capture what is often observed in practice. And our comparison across different combinations of information structure and communication protocol is also confined within this family of symmetric equilibria.

Depending on the price and the payoff functions, not all top strategy profiles will be equilibria. It is easy to check that the values of k for which the top-k strategy profile is an equilibrium will take the form of an interval of consecutive integers. Among these values of k, let k_S denote the sender's favorite top equilibrium and let k_R denote the receiver's favorite top equilibrium. Clearly, k_S will simply be the smallest value of k, so $k_S \leq k_R$. It can be shown that the top-k_R equilibrium is the receiver's most-preferred equilibrium among all possible equilibria (see Appendix 1.7.4).

The top-k_S equilibrium is by definition the sender's most-preferred equilibrium among all *top equilibria*, but it is not necessarily the sender's most-preferred equilibrium among *all* equilibria. In other words, there are other, less natural equilibria, that can increase the buying probability further, and therefore make the sender even better off. Indeed, we can give an example of an *asymmetric* equilibrium that strictly increases the buying probability above the buying probability under the top-k_S equilibrium (see Appendix 1.7.4).

Then one natural question to ask is whether the sender-optimal *symmetric* equilibrium is a top equilibrium. However, this is a hard question to answer at the moment and we leave it to future work. On one hand, there are other symmetric equilibria (we include the construction in Appendix 1.7.4), but we have not found another symmetric equilibrium that increases the buying probability above the sender-preferred top equilibrium. On the other hand, at the current stage we are not able to construct a formal proof to show that such an increase is impossible.

1.5.2 Costly Verification

In the main model we have assumed the verification is costless within the receiver's checking capacity but becomes prohibitively costly beyond this capacity to capture the receiver's time or cognitive constraint as in Glazer and Rubinstein (2004). However, we can relax this assumption and allow for a sufficiently small verification cost and our top-k equilibria can still be sustained. We illustrate this by the following simple setting.

Consider $N = 2$, $v(\theta) = |\theta|$, $|\theta| \sim \text{Binomial}(N, q)$ where $q = 1/2$ and $P = 1$. In the ex ante stage without any verification, under the prior the receiver is indifferent between buying and not buying (note that $\mathbb{E}(|\theta|) = P$) and assume he will not buy which yields payoff 0. Suppose the receiver will incur a cost c for each attribute he checks. Also note that if the receiver checks both attributes, he will not buy since the expected payoff from buying is $\mathbb{E}(|\theta|) - 2c - P < 0$.

To see we still have the top-1 equilibrium, let m_1 denote the message "Attribute 1 is my highest attribute" and m_2 denote the message "Attribute 2 is my highest attribute." In the top-1 equilibrium, the sender with type $(1, 0)$ sends m_1 with probability 1, type $(0, 1)$ sends m_2 with probability 1 and the type $(0, 0)$ and $(1, 1)$ uniformly randomize over these two messages while the receiver checks the recommended attribute and buys if and only if it is good. Clearly, the sender is incentive compatible since she already maximizes the probability of buying. To see the receiver's strategy is a best response, note that upon receiving m_1, the receiver's expected payoff from checking attribute 1 is

$$\frac{1}{4} \times 0 + \frac{3}{4}(1 + \frac{1}{3} \times 1 - P) - c = 1 - \frac{3}{4}P - c,$$

and his expected payoff from checking attribute 2 is

$$\frac{3}{4} \times 0 + \frac{1}{4}(1 + 1 - P) - c = \frac{1}{2} - \frac{1}{4}P - c.$$

When $P = 1$, these two expected payoffs are both equal to $1/4 - c$. Therefore, if $c < 1/4$, the receiver strictly prefers checking exactly one attribute over checking no attribute or both attributes. Similar argument for the case where the receiver obtains m_2. Hence, if the cost of verification is sufficiently small ($c < 1/4$ in this setting), top-1 equilibrium still exists.

In general, when the verification is costly, new issues will arise. For example, whether the receiver checks at all and if he does how many attributes he checks may also depend on the cost, etc. Clearly, this makes the analysis more involved and we leave it to future work. Here, we would like to mention several papers that touch upon related issues. Ben-Porath et al. (2014) characterize a favored-agent mechanism to allocate an indivisible good among a group of senders where monetary transfers are not allowed and the receiver can learn each sender's type at a given cost. In a similar model, Erlanson and Kleiner (2017) study optimal mechanism for the principal in collective choice problems. They show that this mechanism can be implemented as a weighted majority voting rule. In contrast, Mylovanov and Zapechelnyuk (2017) study the allocation of an indivisible prize among multiple agents in a setting where the principal learns the true value from allocating the prize ex post, namely, after the allocation decision has been made.

1.5.3 Equilibrium Refinement

As we have already seen, the top equilibria are not even unique among the symmetric equilibria, although we still think the top equilibria are the most natural equilibria. Therefore, we suffer from a plethora of equilibria as other cheap talk games.

As for further equilibrium refinements, since θ has full support, each node at which the sender selects a message is on the equilibrium path. Since for all messages $A \in \mathcal{M}$, we have $m(A|\theta) > 0$ for some θ, it follows that each information set at which the receiver chooses a set of attributes B to check is on the equilibrium path. The subtlety arises at the information sets indexed by (y, B, A) where the receiver chooses whether to buy the good. Such a node may be off the equilibrium path either because $c(B|A) = 0$ and hence it is inconsistent with the receiver's strategy, or because $m(A|\theta) = 0$ for all θ such that $\theta_B = y$, and hence it is inconsistent with the sender's strategy. As long as a node is consistent with the sender's strategy, there is a natural way to update the receiver's beliefs if he totally mixes. So it is natural to impose either sequential equilibrium or trembling-hand perfect equilibrium. However, even these refinements will not place any restrictions on the receiver's beliefs at nodes where the sender has trembled.

To be clear on this, let's see an example. Suppose $N = 3$ and $n = 2$. For simplicity, suppose the sender is playing the top-1 strategy. Suppose the realized type is $(0, 1, 0)$

and the sender deviates by pointing to attribute 1. At first, the receiver does not know that the sender has deviated because he believes that attribute 1 could be one of the highest attributes. Suppose the receiver's strategy prescribes that, following this message (i.e., "Attribute 1 is my highest attribute"), he checks the recommended attribute 1 and also the unrecommended attribute 2 (recall $n = 2$). Upon seeing that the first attribute is 0 and the second attribute is 1, the receiver now knows that the sender deviated. The receiver also knows the realizations of the first two attributes, but he must form a belief about the third. To form his belief, as a rational Bayesian, the receiver compares the probabilities of (i) $\theta = (0, 1, 0)$ and the sender deviated by pointing to attribute 1; and (ii) $\theta = (0, 1, 1)$ and the sender deviated by pointing to attribute 1. The relative likelihood of these two events is not pinned down because either of them is a "tremble." To pin down the belief, we have to take a position on the relative likelihood of different trembles. All we can say for sure is that the receiver knows $|\theta|$ is either 1 or 2, but the relative probability she assigns to 1 and 2 is undetermined. The conditional expectation of $|\theta|$ can be any number in $[1, 2]$.

1.5.4 Real Cheap Talk?

Although there is no intrinsic content of the "cheap talk" message, we take the intuitive interpretation that the sender announces her k highest attributes in a top-k equilibrium. The information disclosure seems verifiable. However, this is not true since the sender's message does not indicate any ordering among those k attributes and the receiver cannot outright observe the true state of the world. That is, in our model, the sender has no evidence to present. This clarifies the difference with the evidence game defined by Hart et al. (2017) where each type of the sender is characterized by a set of verifiable statements from which the sender chooses. As different messages are available for different types of the sender, the messages amount to evidence. Nevertheless, in our model all the messages are available for all types of the sender and hence the messages are "cheap talk."

1.6 Conclusion

As new technologies revolutionize the collection, storage, and analysis of data, information is becoming an increasingly important commodity. In this paper, we study the scope for persuasion in a static environment where a sender who is perfectly informed about the state costlessly transmits a message to a receiver, who then chooses which aspect of the state to verify. Specifically, we consider a sender-receiver game with a multi-dimensional state. The sender observes the state and costlessly sends a message to the receiver, who selects *some* components of the state to check and then chooses a binary action. Even when the sender always prefers one action independent of the state, we show that for a range of state-dependent preferences for the receiver, there exists a natural family of informative equilibria. In these equilibria, the sender indicates which attributes are highest; the receiver checks some of those attributes and then chooses his action based on their realizations. Across the family of equilibria we construct, the receiver faces a trade-off between the frequency of seeing a good signal and the strength of that signal: a good signal is not representative of good quality when it is observed too often.

We find that compared to alternative communication structures, in the equilibria we characterize even though the sender has state-independent preferences, she can strictly benefit from the ability to communicate. If the sender has commitment power, as in Kamenica and Gentzkow (2011), then she can further increase her utility by committing to randomize between various messages. This is not so surprising by the Bayesian persuasion literature. However, we find that in some non-generic cases, the sender cannot benefit from commitment any more and costless message (cheap talk) can do as well as committing to a signal structure (Bayesian persuasion). Finally, we observe that the receiver's ability to partially verify the state has an ambiguous effect on the sender's utility unless the sender has commitment power, in which case verification can only restrict the set of posteriors the sender can induce.

1.7 Appendix

1.7.1 Proofs for Section 1.3

Theorem 1.1. *Let $k \in [N]$. There exists a top-k equilibrium if and only if*

$$P \leq \mathbb{E}[v(\theta) \mid (\bar{S}_1^k, \underline{S}_1^{N-k}) = (1,0)].$$

Proof. Let $n = 1$ and apply Theorem 1.2. □

Proposition 1.1. *Let $v(\theta) = |\theta|$. There exists a top-1 equilibrium if and only if*

$$P \leq \tfrac{N\mathbb{E}(|\theta|) - \mathbb{E}(|\theta|^2)}{N(1-\pi_0) - \mathbb{E}(|\theta|)}.$$

Proof. Let $v(\theta) = |\theta|$ and $n = 1$. By routine calculation, we have $\mathbb{E}[v(\theta) \mid (\bar{S}_1^k, \underline{S}_1^{N-k}) = (1,0)] = \frac{N\mathbb{E}(|\theta|) - \mathbb{E}(|\theta|^2)}{N(1-\pi_0) - \mathbb{E}(|\theta|)}$. Apply Theorem 1.1. □

Proposition 1.2. *The price threshold $\bar{P}_k$ is strictly increasing in k.*

Proof. Let $n = 1$ and apply Proposition 1.6. □

Proposition 1.3. *The equilibrium buying probability $\mathbb{E}[Z_k]$ is decreasing in k but the conditional expected utility of the receiver $\mathbb{E}[v(\theta) \mid Z_k = 1]$ is increasing in k.*

Proof. For the first statement, note that

$$\mathbb{E}[Z_k] = \frac{\mathbb{E}[k \wedge |\theta|]}{k} = \sum_{i=1}^{k} \pi_{_i} \cdot \frac{i}{k} + \pi_{_{k+1}} + \cdots + \pi_{_N}.$$

Then we have

$$\begin{aligned}
\mathbb{E}[Z_{k+1}] &= \sum_{i=1}^{k+1} \pi_{_i} \cdot \frac{i}{k+1} + \pi_{_{k+2}} + \cdots + \pi_{_N} \\
&= \sum_{i=1}^{k} \pi_{_i} \cdot \frac{i}{k+1} + \pi_{_{k+1}} + \pi_{_{k+2}} + \cdots + \pi_{_N} \\
&< \sum_{i=1}^{k} \pi_{_i} \cdot \frac{i}{k} + \pi_{_{k+1}} + \cdots + \pi_{_N} = \mathbb{E}[Z_k].
\end{aligned}$$

The inequality holds for any $k \in \{1, \ldots, N\}$. Therefore, $\mathbb{E}[Z_k]$ is decreasing in k.

For the second statement, let $\tilde{v} : \{0, \dots, N\} \to \mathbb{R}$ such that $\tilde{v}(i) = v(\theta)$ if $|\theta| = i$.

$$\begin{aligned}
\mathbb{E}[v(\theta)|Z_k = 1] = \frac{\mathbb{E}[v(\theta)\mathbb{1}(Z_k = 1)]}{\Pr(Z_k = 1)} &= \frac{\sum_{i=1}^{N} \tilde{v}(i) \cdot \Pr(|\theta| = i) \cdot \Pr(Z_k = 1 \mid |\theta| = i)}{\mathbb{E}[Z_k]} \\
&= \frac{\sum_{i=1}^{k} \tilde{v}(i) \cdot \pi_i \cdot \frac{i}{k} + \sum_{i=k+1}^{N} \tilde{v}(i) \cdot \pi_i \cdot 1}{\sum_{i=1}^{k} \pi_i \cdot \frac{i}{k} + \pi_{k+1} + \cdots + \pi_N} \\
&= \frac{\frac{k}{k+1} \sum_{i=1}^{k} \tilde{v}(i) \cdot \pi_i \cdot \frac{i}{k} + \frac{k}{k+1} \sum_{i=k+1}^{N} \tilde{v}(i) \cdot \pi_i}{\frac{k}{k+1} \sum_{i=1}^{k} \pi_i \cdot \frac{i}{k} + \frac{k}{k+1} \sum_{i=k+1}^{N} \pi_i}.
\end{aligned}$$

Then we have

$$\begin{aligned}
\mathbb{E}[v(\theta)|Z_{k+1} = 1] &= \frac{\sum_{i=1}^{k+1} \tilde{v}(i) \cdot \pi_i \cdot \frac{i}{k+1} + \sum_{i=k+2}^{N} \tilde{v}(i) \cdot \pi_i \cdot 1}{\mathbb{E}[Z_{k+1}]} \\
&= \frac{\sum_{i=1}^{k} \tilde{v}(i) \cdot \pi_i \cdot \frac{i}{k+1} + \sum_{i=k+1}^{N} \tilde{v}(i) \cdot \pi_i}{\sum_{i=1}^{k} \pi_i \cdot \frac{i}{k+1} + \pi_{k+1} + \pi_{k+2} + \cdots + \pi_N} \\
&= \frac{\frac{k}{k+1} \sum_{i=1}^{k} \tilde{v}(i) \cdot \pi_i \cdot \frac{i}{k} + \frac{k}{k+1} \sum_{i=k+1}^{N} \tilde{v}(i) \cdot \pi_i + \frac{1}{k+1} \sum_{i=k+1}^{N} \tilde{v}(i) \cdot \pi_i}{\frac{k}{k+1} \sum_{i=1}^{k} \pi_i \cdot \frac{i}{k} + \frac{k}{k+1} \sum_{i=k+1}^{N} \pi_i + \frac{1}{k+1} \sum_{i=k+1}^{N} \pi_i}.
\end{aligned}$$

Let

$$a := \frac{k}{k+1} \sum_{i=1}^{k} \tilde{v}(i) \cdot \pi_i \cdot \frac{i}{k} + \frac{k}{k+1} \sum_{i=k+1}^{N} \tilde{v}(i) \cdot \pi_i, \quad b := \frac{k}{k+1} \sum_{i=1}^{k} \pi_i \cdot \frac{i}{k} + \frac{k}{k+1} \sum_{i=k+1}^{N} \pi_i,$$
$$c := \frac{1}{k+1} \sum_{i=k+1}^{N} \tilde{v}(i) \cdot \pi_i, \quad d := \frac{1}{k+1} \sum_{i=k+1}^{N} \pi_i,$$
$$e := \frac{k}{k+1} \sum_{i=1}^{k} \tilde{v}(i) \cdot \pi_i \cdot \frac{i}{k}, \quad f := \frac{k}{k+1} \sum_{i=1}^{k} \pi_i \cdot \frac{i}{k}.$$

Then,

$$a = e + k \cdot c, \quad b = f + k \cdot d,$$
$$\mathbb{E}[v(\theta)|Z_k = 1] = \frac{a}{b}, \quad \mathbb{E}[v(\theta)|Z_{k+1} = 1] = \frac{a+c}{b+d}.$$

So we have $a \cdot (f + k \cdot d) = b \cdot (e + k \cdot c) \Leftrightarrow (ad - bc)k = be - af$.
Note that

$$\begin{aligned}
be - af = &\Big(\frac{k}{k+1}\sum_{i=1}^{k} \pi_{_i} \cdot \frac{i}{k} + \frac{k}{k+1}\sum_{i=k+1}^{N} \pi_i\Big)\Big(\frac{k}{k+1}\sum_{i=1}^{k} \tilde{v}(i) \cdot \pi_{_i} \cdot \frac{i}{k}\Big) - \\
&\Big(\frac{k}{k+1}\sum_{i=1}^{k} \tilde{v}(i) \cdot \pi_{_i} \cdot \frac{i}{k} + \frac{k}{k+1}\sum_{i=k+1}^{N} \tilde{v}(i) \cdot \pi_{_i}\Big)\Big(\frac{k}{k+1}\sum_{i=1}^{k} \pi_{_i} \cdot \frac{i}{k}\Big) \\
&= (\frac{k}{k+1})^2\Big[\Big(\sum_{i=k+1}^{N} \pi_i\Big)\Big(\sum_{i=1}^{k} \tilde{v}(i) \cdot \pi_{_i} \cdot \frac{i}{k}\Big) - \Big(\sum_{i=k+1}^{N} \tilde{v}(i) \cdot \pi_{_i}\Big)\Big(\sum_{i=1}^{k} \pi_{_i} \cdot \frac{i}{k}\Big)\Big] \\
&< (\frac{k}{k+1})^2\Big[\Big(\sum_{i=k+1}^{N} \pi_i\Big)\Big(\sum_{i=1}^{k} \tilde{v}(k+1) \cdot \pi_{_i} \cdot \frac{i}{k}\Big) - \Big(\sum_{i=k+1}^{N} \tilde{v}(i) \cdot \pi_{_i}\Big)\Big(\sum_{i=1}^{k} \pi_{_i} \cdot \frac{i}{k}\Big)\Big] \\
&= (\frac{k}{k+1})^2\Big[\Big(\sum_{i=k+1}^{N} \tilde{v}(k+1) \cdot \pi_i\Big)\Big(\sum_{i=1}^{k} \pi_{_i} \cdot \frac{i}{k}\Big) - \Big(\sum_{i=k+1}^{N} \tilde{v}(i) \cdot \pi_{_i}\Big)\Big(\sum_{i=1}^{k} \pi_{_i} \cdot \frac{i}{k}\Big)\Big] \\
&< 0.
\end{aligned}$$

The first inequality follows from the fact that $\tilde{v}(k+1) > \tilde{v}(i)$ for all $i \leq k$ and the last inequality follows from the fact that $\tilde{v}(k+1) < \tilde{v}(i)$ for all $i \geq k+2$. Therefore, we have

$$(ad - bc)k = be - af < 0 \Rightarrow ad < bc \Rightarrow \frac{a}{b} < \frac{a+c}{b+d} \Leftrightarrow \mathbb{E}[v(\theta)|Z_k = 1] < \mathbb{E}[v(\theta)|Z_{k+1} = 1].$$

The inequality holds for any $k \in \{1, \dots, N\}$. Therefore, $\mathbb{E}[v(\theta)|Z_k = 1]$ is increasing in k.

□

Proposition 1.4. *The ranking of trading probability in the sender-optimal equilibrium of different settings is as follows:*

$$NC \leq CT \leq BP \text{ and } NCV \leq CTV \leq BPV \leq BP.$$

Proof. First, we show NC≤CT≤BP. If there is no verification, as we discussed in the main text, the only equilibrium in the cheap talk game is the babbling equilibrium in which the receiver ignores the sender's message and makes a purchase decision based on his prior belief about the sender's type. Given the receiver's strategy, the sender finds it not worthwhile to convey any information in the message she sends. This is due to the sender's state-independent preferences in the sense that she only aims to maximize the probability of buying. If there are two messages one of which induces a higher probability of buying, then all types of the sender will be pooling on that message. Therefore, cheap talk (CT) is equivalent to no communication (NC) between the two parties. And the probability of buying in NC is equal to that in CT.

If the sender can commit to a statistical experiment (BP), one option for the sender which is always available is an experiment that generates a single signal about the sender's type, i.e., there is actually no experiment. The receiver then makes a purchase decision after observing this signal. Hence, the probability of buying is determined by the two parties' common prior about the sender's type. Thus, the probability of buying induced by this experiment is equal to that in NC or CT. Clearly, the seller can be weakly better off by designing a non-degenerate experiment. Therefore, the probability of buying in BP is weakly higher than NC or CT.

Second, we show NCV≤CTV≤BPV≤BP. We start with showing NCV≤CTV. Clearly, $\mathbb{E}[v(\theta) \mid \theta_i = 0] \leq \mathbb{E}[v(\theta)] \leq \mathbb{E}[v(\theta) \mid \theta_i = 1]$ for any $i \in \{1, \ldots, N\}$. If $P \leq \mathbb{E}[v(\theta) \mid \theta_i = 0]$, NCV is trivially equal to CTV since the price is too low so that the receiver always buys. Then consider $\mathbb{E}[v(\theta) \mid \theta_i = 0] < P$. We further restrict attention to prices that can sustain top-k equilibria, i.e., $P \leq \bar{P}_k$. Then (i) if $\mathbb{E}[v(\theta) \mid \theta_i = 1] < P \leq \bar{P}_k$, then the probability of buying in NCV is 0 but positive in CTV; (ii) if $P \leq \bar{P}_k \leq \mathbb{E}[v(\theta) \mid \theta_i = 1]$ or $P \leq \mathbb{E}[v(\theta) \mid \theta_i = 1] \leq \bar{P}_k$, then the receiver will buy only if she sees a "1" in NCV and the corresponding probability of buying is $\mathbb{E}(|\theta|)/N$. Recall that in CTV, the probability of buying of a top-k equilibrium is

$$\begin{aligned}
\mathbb{E}[Z_k] &= \sum_{i=1}^{k} \pi_i \cdot \frac{i}{k} + \sum_{i=k+1}^{N} \pi_i \\
&> \sum_{i=1}^{k} \pi_i \cdot \frac{i}{N} + \sum_{i=k+1}^{N} \pi_i \cdot \frac{i}{N} = \frac{E(|\theta|)}{N}.
\end{aligned}$$

Next, we show CTV≤BPV≤BP. If the sender can commit to a signal structure, let $b^{\mathrm{BP}} = (b_0, \ldots, b_N)$ be the buying vector under the optimal persuasion mechanism, where b_j denotes the probability of buying conditional on $|\theta|= j$. It can be computed by maximizing the sender's utility subject to the receiver's utility constraint. And it will be of the form $b^{\mathrm{BP}} = (0, \ldots, 0, \alpha, 1, \ldots, 1)$, where $\alpha \in (0, 1]$. Then the optimal persuasion mechanism must be split into two cases. If $b_1^{\mathrm{BP}} = 1$, then the sender can simply tell the receiver whether to buy or not. If $b_1^{\mathrm{BP}} < 1$, however, this strategy will not work, and the sender must instead point to an attribute and tell the receiver to buy if and only this attribute is one. The attribute should be chosen so that it is distributed uniformly and this behavior results in the desired buying vector.

Why this difference? The key is that the sender must find a way to achieve his optimal buying vector without giving the receiver any valuable information. If $b_1^{\mathrm{BP}} < 1$, then the optimal buying vector does not always require buying when there is some good attribute.

This just comes down to whether the top-1 strategy profile is an equilibrium. If it is, then under the optimal persuasion mechanism, the sender will always recommend buying when there is at least one attribute. If the top-1 strategy profile is not an equilibrium, then the sender will not always recommend buying even when there is some positive attributes. In this case, telling the agent to buy could lead into trouble because the receiver may be better off deviating and checking a random attribute, secure in the knowledge that he has dodged some very bad states.

In other words, with the extra instrument of commitment, the sender can always replicate the sender-optimal top equilibrium in CTV, and by manipulating her randomization device, the sender can achieve a (weakly) better buying vector without giving the receiver any more valuable information.

□

1.7.2 Proofs for Section 1.4

Theorem 1.2. *Suppose $n \leq N - 2$. For each $k = 1, \ldots, N$, there exists a top-k equilibrium if and only if*

$$P \leq \mathbb{E}[v(\theta) \mid (\bar{S}_n^k, \underline{S}_n^{N-k}) = (n \wedge k, 0)].$$

If $n = N - 1$, it goes through except the "only if" statement holds for all but finitely *many prices.*[16]

To make the exposition more clear, we relabel the state θ as X and therefore $v(\theta) = v(|X|)$, and denote a specific realization of the random variable X by x.

In order to prove Theorem 1.2, a number of statistical lemmas are needed. To simplify the arguments below, we first construct a random variable as follows. Given a message strategy m, define an $\mathcal{M}$-valued random variable $M = M(m)$ on the same probability space as X, with the joint distribution of (X, M) determined by π and m. That is, the joint probability mass function is given by $f(x, A) = m(A|x)\pi(x)$. We will refer to M as the *random message induced by* m. Note, however, that the sender's strategy is still a function.

Next, we introduce some notations for the hypergeometric distribution. Let $\mathrm{HG}(n, K, N)$ denote the hypergeometric distribution when the sample size is n, the number of successes is K, and the total population size is N. Denote the corresponding probability mass function by

$$p(x; n, K, N) = \frac{\binom{K}{x}\binom{N-K}{n-x}}{\binom{N}{n}}, \quad 1 \leq x \leq n \wedge K. \tag{1.1}$$

Lemma 1.1. Let M be the random message induced by the strategy m^k. Conditional on M and $|X|$,

(a) The random vectors X_M and $X_{[N]\setminus M}$ are independent.

(b) The components X_i for $i \in M$ are exchangeable, and for any $I \subset M$, we have $|X_I| \sim \mathrm{HG}(|I|, k \wedge |X|, k)$.

(c) The components X_j for $j \in [N] \setminus M$ are exchangeable, and for any $J \subset [N] \setminus M$, we have $|X_J| \sim \mathrm{HG}(|J|, |X| - k \wedge |X|, N - k)$.

Proof. By exchangeability, $\pi(x) = \pi(x')$ whenever $|x| = |x'|$. By construction, for any $A \in \mathcal{P}_k$,

$$\Pr((X, M) = (x, A)) = \begin{cases} \pi(x)/\binom{N-|x|}{k-|x|} & \text{if } |x_A| = k \wedge |x| = |x|, \\ \pi(x)/\binom{|x|}{k} & \text{if } |x_A| = k \wedge |x| = k, \\ 0 & \text{otherwise.} \end{cases}$$

[16]Specifically, when $n = N - 1$ and $k \leq N - 2$, we can also construct top-k equilibrium if $P = \mathbb{E}[v(\theta) \mid (\bar{S}_n^k, \underline{S}_n^{N-k}) = (n \wedge k, j)]$ for some $j = 1, \ldots, N - k - 1$.

Notice that this expression depends only on $|x|$ and $|x_A|$. Therefore, conditional on the event that $|X|= s$ and $M = A$, the vector $(X_A, X_{[N]\setminus A})$ is uniformly distributed over the set

$$\{x \in \{0,1\}^A \mid |x|= k \wedge s\} \times \{y \in \{0,1\}^{[N]\setminus A} \mid |y|= s - k \wedge s\}.$$

This proves (a), and parts (b) and (c) now follow from standard computations, which are omitted. □

Next, we define random variables that are equal in distribution to $|X_I|$ and $|X_J|$ from Lemma 1.1. For $0 \le a \le b \le N$, define the random variables $\bar{S}_a^b$ and $\underline{S}_a^b$ on the same probability space as (X, M) as follows. Conditional on X, these random variables are mutually independent of each other and M, with

$$\bar{S}_a^b \sim \mathrm{HG}(a, b \wedge |X|, b) \quad \text{and} \quad \underline{S}_a^b \sim \mathrm{HG}(a, |X|-(N-b) \wedge |X|, b).$$

The idea is that $\bar{S}_a^b$ ($\underline{S}_a^b$) is equal in distribution to the sum of a attributes chosen uniformly from b of the top (bottom) attributes.[17] To simplify the notations below, set $\bar{S}_a^b = \bar{S}_b^b$ and $\underline{S}_a^b = \underline{S}_b^b$ for $a > b$.

After checking the attributes in some set B, the receiver's information is summarized by the pair (X_B, M). We introduce a pair of simpler statistics. Let

$$T(X_B, M) = (|X_{B\cap M}|, |X_{B\setminus M}|) \quad \text{and} \quad U(M) = (|B \cap M|, |B \setminus M|). \tag{1.2}$$

In particular, T and U depend on the set B of attributes that the receiver checks. Note that the symbol $|\cdot|$ denotes the sum of components of vectors in T and the cardinality of sets in U.

Lemma 1.2. Let M be the random message induced by the strategy m^k. Fix a nonempty subset B of $[N]$. Given the sample (X_B, M), the statistic (T, U) defined in (1.2) is sufficient for $|X|$.

Proof. With p denoting the hypergeometric probability mass function (see (1.1)),

[17] To determine these top and bottom attributes, ties are broken uniformly.

Lemma 1.1 gives

$$\begin{aligned}
&\Pr((X_B, M) = (x_B, A) \mid |X| = s) \\
&= \binom{s}{k \wedge s}^{-1} p(|x_{B\cap A}|; |B \cap A|, k \wedge s, k) p(|x_{B\setminus A}|; |B \setminus A|, s - k \wedge s, N - k).
\end{aligned}$$

Since this expression depends on (x_B, A) only through

$$(T, U) = (|x_{B\cap A}|, |x_{B\setminus A}|, |B \cap A|, |B \setminus A|),$$

sufficiency follows. □

This statistic (T, U) satisfies an intuitive and useful monotonicity property. For any $u = (u_1, u_2) \in \operatorname{supp} U$, define the conditional support of T, given $U = u$, by

$$\mathbf{S}(u) = \{t \mid (t, u) \in \operatorname{supp}(T, U)\}.$$

Note that

$$\mathbf{S}(u) = \{(0, 0), (1, 0), \ldots, (u_1, 0), (u_1, 1) \ldots, (u_1, u_2)\},$$

which is obviously totally ordered by the product order on $\mathbb{Z}^2$. We will use the usual symbol $\geq$ to denote the product order; the meaning should be clear from the context.[18]

We are interested in whether one signal realization is "more favorable" than another in the spirit of Milgrom (1981).

Definition. Given realizations (t, u) and (t', u') in $\operatorname{supp}(T, U)$, we say (t, u) is *more favorable* than (t', u'), denoted $(t, u) \succ_{\text{fav}} (t', u')$, if the conditional distribution of $|X|$ given $(T, U) = (t, u)$ first-order stochastically dominates the conditional distribution of $|X|$ given $(T, U) = (t', u')$, for any strictly positive distribution of $|X|$.[19]

Lemma 1.3. Fix a nonempty subset B of $[N]$. For any $(t, u), (t', u) \in \operatorname{supp}(T, U)$, $t > t'$ implies $(t, u) \succ_{\text{fav}} (t', u)$.

Proof. Fix $(t, u), (t', u) \in \operatorname{supp}(T, U)$ with $t > t'$. Let $f(\cdot|\cdot)$ denote the conditional probability mass function of (T, U) given $|X|$. Following the argument of Milgrom

[18] Recall that the product order means $(a, b) \geq (a', b')$ if and only if $a \geq a'$ and $b \geq b'$.

[19] Milgrom (1981) defines a weak and a stronger notion of favorability. Our definition lies between the two in strength.

(1981), it suffices to show that

$$f(t,u|s)f(t',u|s') \geq f(t,u|s')f(t',u|s),$$

for all $s > s'$ with strict inequality for some $s > s'$.[20] Divide each side by $f(u|s)f(u|s') > 0$ and then plug in the expressions for f to obtain

$$\begin{aligned}
&p(t_1; u_1, k \wedge s, k)p(t_1'; u_1, k \wedge s', k) \\
&\qquad \cdot p(t_2; u_2, s - k \wedge s, N-k)p(t_2'; u_2, s' - k \wedge s', N-k) \\
&\geq p(t_1'; u_1, k \wedge s, k)p(t_1; u_1, k \wedge s', k) \\
&\qquad \cdot p(t_2'; u_2, s - k \wedge s, N-k)p(t_2; u_2, s' - k \wedge s', N-k).
\end{aligned}$$

We will use the fact that for the hypergeometric family of distributions (parametrized by the number of successes) has a strictly monotone likelihood ratio. Suppose $s > s'$. Then $k \wedge s \geq k \wedge s'$ with strict inequality if $s' < k$; likewise, $s - k \wedge s \geq s' - k \wedge s'$ with strict inequality if $s > k$. Comparing the first two terms on each side of the inequality and then the last two terms on each side gives the weak inequality. We claim that the inequality is strict if $s = t_1 + t_2$ and $s' = t_1' + t_2'$, which satisfies $s > s'$ since $t > t'$. Clearly, all terms on the LHS are strictly positive. We separate into cases. If $t_1 > t_1'$ then $t_1' < u_1 \wedge k$ and hence $t_2' = 0$ so $s' < k$. Therefore, the product of the first two terms is strictly greater on the left than on the right. If $t_2 > t_2'$, then $t_2 > 0$ so $s > k$. Therefore, the product of the last two terms is strictly greater on the left than on the right. In either case, the inequality is strict. □

With these statistical lemmas established, we now define the receiver's strategy as follows. If he can check all the attributes recommended by the sender, he does so (and does not check any others). If he cannot check all the recommended attributes, then he checks n of them. Then he buys if the number of ones he sees exceeds some threshold. Formally, let $c^*(A)$ be the uniform distribution over $\mathcal{P}_{n \wedge k}(A)$, the family of $(n \wedge k)$-element subsets of A, and let $b^*(y, B, A) = \mathbb{1}\{|y| \geq \bar{s}^*\}$ where the threshold $\bar{s}^*$ is defined by

$$\bar{s}^*(P,k,n) = (n \wedge k + 1) \wedge \min\Big\{s \in \{0, \ldots, n \wedge k\} \mid \mathbb{E}[v(|X|) \mid \bar{S}_n^k = s] \geq P\Big\}.$$

[20] Milgrom (1981) uses a slightly different notion of favorability, but his argument can be easily modified to apply to our definitions.

Notice that under this strategy, the receiver buys the product when indifferent. We also define the analogous threshold $\underline{s}^*$ for unrecommended attributes by

$$\underline{s}^*(P,k,n) = (n \wedge (N-k)+1) \wedge \min\Big\{ s \in \{0,\ldots,n\wedge(N-k)\} \mid \mathbb{E}[v(|X|) \mid \underline{S}_n^{N-k} = s] \geq P \Big\}.$$

Now we turn to the proof proper.

Proof. First suppose the price inequality in the theorem statement holds, i.e., $P \leq \mathbb{E}[v(|X|) \mid (\bar{S}_n^k, \underline{S}_n^{N-k}) = (n \wedge k, 0)]$. We show that (m^k, c^*, b^*) is a top-k equilibrium. Suppose counterfactually that the receiver were allowed to check $n \wedge k$ recommended attributes and also $n \wedge (N-k)$ unrecommended attributes. By Lemmas 1.2 and 1.3, there exists a best response for the receiver that is a threshold strategy in T, i.e., buy if and only if $T \geq t^*$ for some $t^* \in \mathbf{S}(n \wedge k, n \wedge (N-k))$. Since $P \leq \mathbb{E}[v(|X|) \mid (\bar{S}_n^k, \underline{S}_n^{N-k}) = (n \wedge k, 0)]$, the threshold t^* can be chosen so that $t^* \geq (n \wedge k, 0)$. Therefore the threshold strategy is measurable with respect to $n \wedge k$ of the recommended attributes, and hence (c^*, b^*) achieves the same payoff as the best response in the counterfactual game. We conclude that (c^*, b^*) remains a best response in the actual game in which the receiver is only allowed to check n attributes in total, and that the receiver is playing a best response to sender's strategy.

Given the receiver's strategy (c^*, b^*), the sender's strategy is clearly a best response because it maximizes the probability of purchase. Finally, (m^k, c^*, b^*) is a top-k equilibrium because, in particular, the receiver buys with probability one whenever $|X| \geq n \wedge k$, which has positive probability since $n < N$ and X has full support.

Next, suppose the price inequality in the theorem statement does not hold. Suppose for a contradiction that there is a top-k equilibrium (m, c, b) with $m = m^k$. First assume

$$P \neq \mathbb{E}[v(|X|) \mid (\bar{S}_n^k, \underline{S}_n^{N-k}) = (n \wedge k, j)] \tag{1.3}$$

for all $j = 1, \ldots, n \wedge (N-k)$. Following the same reasoning as above, it can be shown that under any best response, the receiver must check $n \wedge (N-k)$ of the unrecommended attributes and must buy if and only if $t_2 \geq t_2^*$ for some $t_2^* > 1$. Since this is a top-k equilibrium, we must have $t_2^* \leq n \wedge (N-k)$. But then the sender can profitably deviate when $|X| = k + \bar{s}^* - 1$. If the sender follows the top-k strategy and chooses A with $|X_A| = k$, then $|X_{[N]\setminus B}| = \bar{s}^* - 1$, and the receiver will never buy the good. If instead, the sender sends a message A such that $|X_A| = k - 1$, then

$|X_{[N]\setminus A}|= \bar{s}^*$ and with positive probability, $|X_{B\setminus A}|= \bar{s}^*$ and the receiver buys. This contradiction completes the proof under the generic assumption (1.3).

Lastly, suppose

$$P = \mathbb{E}[v(|X|) \mid (\bar{S}_n^k, \underline{S}_n^{N-k}) = (n \wedge k, \hat{j})]$$

for some $\hat{j} \geq 1$. Now the receiver's optimal strategy is not uniquely pinned down because the buying probability is completely free when he sees j. However, the argument still goes through as before. If $n \geq N-k$, there is a special case to consider because there may exist (A, B) on the equilibrium path with $|B \setminus A|= N - k - 1$. Then the receiver's strategy is pinned down. He must buy if and only if $|y| \geq \underline{s}^* - 1$, at least along the equilibrium path. Suppose $|X|= k + s^* - 2$ and $|X_A|= k - 1$ and $X_i = 0$ for the unique attribute $i \notin A \cup B$. On the equilibrium path, the receiver will never buy. By sending message A instead, the receiver will buy with positive probability. This completes the proof.

□

Theorem 1.2 still holds if the receiver can check the attributes sequentially rather than simultaneously.

Proof. Clearly if the inequality is satisfied, the buyer is doing as well as she can. If it is violated, sequential checking must have the same result in the particular cases shown, and hence there is no benefit. □

Proposition 1.5. *Both thresholds $\bar{s}^*(P, k, n)$ and $\underline{s}^*(P, k, n)$ are weakly increasing in P and n but weakly decreasing in k.*

Proof. Recall

$$\bar{s}^*(P, k, n) = (n \wedge k + 1) \wedge \min\left\{s \in \{0, \ldots, n \wedge k\} \mid \mathbb{E}[v(|X|) \mid \bar{S}_n^k = s] \geq P\right\},$$

and

$$\underline{s}^*(P, k, n) = (n \wedge (N-k)+1) \wedge \min\left\{s \in \{0, \ldots, n \wedge (N-k)\} \mid \mathbb{E}[v(|X|) \mid \underline{S}_n^{N-k} = s] \geq P\right\}.$$

As P increases, the minima are taken over smaller sets and hence increase. For the comparative statics in n and k, it suffices to check that both conditional expectations in these definitions are weakly decreasing in n and weakly increasing in k. To simplify notation, let $\bar{f}_{n,k}(\cdot|\cdot)$ denote the conditional probability mass function of $\bar{S}_n^k$ given $|X|$. To show that $\mathbb{E}[v(|X|)|\bar{S}_n^k = s]$ is decreasing in n, it suffices to check that for all t and $s > s'$, we have

$$\bar{f}_{n,k}(t|s)\bar{f}_{n',k}(t|s') \leq \bar{f}_{n,k}(t|s')\bar{f}_{n',k}(t|s)$$

for all k and $n > n'$. To show that $\mathbb{E}[v(|X|) \mid \bar{S}_n^k = s]$ is decreasing in k, it suffices to check that for all t and $s > s'$, we have

$$\bar{f}_{n,k}(t|s)\bar{f}_{n,k'}(t|s') \geq \bar{f}_{n,k}(t|s')\bar{f}_{n,k'}(t|s)$$

for all n and $k > k'$. To establish the comparative statics in $\underline{s}^*(P, k, n)$, we simply prove the equivalent inequalities with $\underline{f}$ in place of $\bar{f}$, where $\underline{f}(\cdot|\cdot)$ denotes the conditional probability mass function of $\underline{S}_n^{N-k}$ given $|X|$. These inequalities can all be verified from the standard properties of the hypergeometric distribution listed in Appendix 1.7.3.

□

Proposition 1.6. *For any fixed n, the price threshold $P^*(k, n)$ is strictly increasing in k.*

Proof. Fix n. For any k, let

$$\bar{f}(k|s) = \Pr(\bar{S}_n^k = n \wedge k \mid |X|= s), \quad \underline{f}(k|s) = \Pr(\underline{S}_n^{N-k} = 0 \mid |X|= s).$$

By conditional independence,

$$\bar{f}(k|s)\underline{f}(k|s) = \Pr((\bar{S}_n^k, \underline{S}_n^{N-k}) = (n \wedge k, 0) \mid |X|= s).$$

We need to show that for all $k > k'$, we have

$$\bar{f}(k|s)\underline{f}(k'|s') \geq \bar{f}(k'|s)\underline{f}(k|s'),$$

for all $s > s'$, with strict inequality for some $s > s'$ (which may depend on k). For strictness simply take $s = k$ and $s' = k'$. Then the LHS is unity but the RHS is strictly less than unity since $\underline{f}(k|s') < 1$. For the weak property, it suffices to check the monotone likelihood ratio property (MLRP) separately for $\bar{f}$ and $\underline{f}$. We begin with $\bar{f}$. By definition,

$$\bar{f}(k|s) = p(n \wedge k; n \wedge k, k \wedge s, k) = \frac{\binom{k\wedge s}{n\wedge k}}{\binom{k}{n\wedge k}}.$$

For convenience, we will work with s and $s-1$ and k and $k-1$. It suffices to prove the result when the terms are positive.

If $s \geq k$, then the result is trivial, so we may assume $s < k$ and and hence $s-1 < k-1$. Hence,

$$\begin{aligned}
\frac{\bar{f}(k|s)\bar{f}(k-1|s-1)}{\bar{f}(k|s-1)\bar{f}(k-1|s)} &= \frac{\binom{s}{n\wedge k}\binom{s-1}{n\wedge(n-1)}}{\binom{s-1}{n\wedge k}\binom{s}{n\wedge(k-1)}} \\
&= \frac{(s-1-n\wedge k)!\,(w-n\wedge(k-1))!}{(s-n\wedge k)!\,(s-1-n\wedge(k-1))!}, \\
&= \frac{s-n\wedge(k-1)}{s-n\wedge k} \\
&> 1.
\end{aligned}$$

Now we give the similar proof for $\underline{f}$. By definition,

$$\underline{f}(k|s) = p(0; n\wedge(N-k), s-k\wedge s, N-k) = \frac{\binom{N-k-(s-k\wedge s)}{n\wedge(N-k)}}{\binom{N-k}{n\wedge(N-k)}}.$$

The result is trivial if $s \leq k$, so we may assume $s > k$. Hence,

$$\begin{aligned}
\frac{\underline{f}(k|s)\underline{f}(k-1|s-1)}{\underline{f}(k|s-1)\underline{f}(k-1|s)} &= \frac{\binom{N-s}{n\wedge(N-k)}\binom{N-s+1}{n\wedge(N-k+1)}}{\binom{N-s+1}{n\wedge(N-k)}\binom{N-s}{n\wedge(N-k+1)}} \\
&= \frac{(N-s+1-n\wedge(N-k))!\,(N-s-n\wedge(N-k+1))!}{(N-s-n\wedge(N-k))!\,(N-s+1-n\wedge(N-k+1))!} \\
&= \frac{N-s+1-n\wedge(N-k)}{N-s+1-n\wedge(N-k+1)} \\
&> 1.
\end{aligned}$$

This completes the proof.

□

1.7.3 Statistical Background

Exchangeability

A vector of random variables $X = (X_i)_{i \in I}$ (with finite index set I) is called exchangeable if

$$(X_i)_{i \in I} =_d (X_{\tau(i)})_{i \in I}$$

for any permutation τ on I.

Lemma 1.4. Suppose $X = (X_i)_{i \in I}$ is exchangeable. Then for any $k \in I$,

$$Cov(X_k, \textstyle\sum_{i \in I} X_i) \geq 0.$$

Proof. By exchangeability, $\mathrm{Cov}(X_k, \sum_{i \in I} X_i)$ is independent of k. Therefore,

$$0 \leq \mathrm{Var}(\sum_{i \in I} X_i) = \sum_{j \in I} \mathrm{Cov}(X_j, \sum_{i \in I} X_i) = |I| \mathrm{Cov}(X_k, \sum_{i \in I} X_i)$$

for any $k \in I$, as needed. □

When the components of X are binary, Lemma 1.4 implies that observing a high realization of one component can only increase the expected value of the sum $\sum_{i \in I} X_i$.

Lemma 1.5. Suppose $X = (X_i)_{i \in I} \in \{0, 1\}^I$ is exchangeable and

$$0 < Pr(X_1 = 1) < 1.$$

Then for any $k \in I$,

$$\mathbb{E}[\sum_{i \in I} X_i \mid X_k = 0] \leq \mathbb{E}[\sum_{i \in I} X_i] \leq \mathbb{E}[\sum_{i \in I} X_i \mid X_k = 1].$$

Proof. By exchangeability, $\Pr(X_k = 1)$ does not depend on k. Denote this common value by q. By assumption, $q \in (0, 1)$. By Lemma 1.4,

$$\begin{aligned} q\mathbb{E}[\sum_{i \in I} X_i] &= \mathbb{E}[X_k] \cdot \mathbb{E}[\sum_{i \in I} X_i] \\ &\leq \mathbb{E}[X_k \cdot \sum_{i \in I} X_i] \\ &= q\mathbb{E}[\sum_{i \in I} X_i \mid X_k = 1], \end{aligned}$$

where the outer equalities hold because $X_k \in \{0,1\}$. Then we have $\mathbb{E}[\sum_{i\in I} X_i] \leq \mathbb{E}[\sum_{i\in I} X_i \mid X_k = 1]$ since $q \in (0,1)$. Therefore,

$$\begin{aligned}\mathbb{E}[\sum_{i\in I} X_i] &= (1-q)\mathbb{E}[\sum_{i\in I} X_i \mid X_k = 0] + q\mathbb{E}[\sum_{i\in I} X_i \mid X_k = 1] \\ &\geq (1-q)\mathbb{E}[\sum_{i\in I} X_i \mid X_k = 0] + q\mathbb{E}[\sum_{i\in I} X_i]\end{aligned}$$

which implies

$$\mathbb{E}[\sum_{i\in I} X_i] - q\mathbb{E}[\sum_{i\in I} X_i] \geq (1-q)\mathbb{E}[\sum_{i\in I} X_i \mid X_k = 0],$$

so

$$(1-q)\mathbb{E}[\sum_{i\in I} X_i] \geq (1-q)\mathbb{E}[\sum_{i\in I} X_i \mid X_k = 0].$$

Then we have $\mathbb{E}[\sum_{i\in I} X_i \mid X_k = 0] \leq \mathbb{E}[\sum_{i\in I} X_i]$ since $q \in (0,1)$. This completes the proof. □

Hypergeometric Distribution

We now check that the hypergeometric distribution satisfies the needed monotone likelihood ratio properties (MLRP). We are interested in the hypergeometric probability mass function

$$p(x; n, K, N) = \frac{\binom{K}{x}\binom{N-K}{n-x}}{\binom{N}{n}}, \quad 1 \leq x \leq n \wedge K.$$

We claim that the hypergeometric distribution satisfies the MLRP with respect to n and K and the reverse MLRP with respect to N. In words, no matter your prior over one of the parameters, observing a larger realization of X will always cause you to update your beliefs in the direction of larger n, larger K, but smaller N. To see this, we will use the shorthand $f(x|n)$ when K and N are to be held fixed, and likewise for $f(x|K)$ and $f(x|N)$.

We begin with K. We need to show that for all $x > x'$ and $K > K'$,

$$f(x|K)f(x'|K') \geq f(x'|K)f(x|K').$$

It suffices to verify this for $x' = x - 1$. Suppose $x' = x - j$. For each $i = 0, \ldots, j-1$, we have

$$f(x-i|K)f(x-i-1|K') \geq f(x-i-1|K)f(x-i|K').$$

Taking the product of these j inequalities gives a telescoping product. All terms except those involving x and $x' = x - j$ cancel and we are left with the desired inequality

$$f(x|K)f(x'|K') \geq f(x'|K)f(x|K').$$

In each case the support shifts in the desired direction so that if $f(x|K) = 0$, then either x is above the support in which case $f(x|K') = 0$, or x is below the support in which case $f(x'|K) = 0$. Therefore, it suffices to verify the desired inequality when all terms are positive. Thus, it suffices to check the following three inequalities:

$$\begin{aligned} \frac{f(x|K)f(x-1|K')}{f(x|K;)f(x-1|K)} &\geq 1, \\ \frac{f(x|n)f(x-1|n')}{f(x|n')f(x-1|n)} &\geq 1, \\ \frac{f(x|N)f(x-1|N')}{f(x|N')f(x-1|N)} &\leq 1. \end{aligned}$$

We now prove these in turn. In each case, we can simplify the ratio by noting that the only terms that do not cancel are the factorials involving both x and the parameter of interest.

For the MLRP in K, we have

$$\begin{aligned} &\frac{f(x|K)f(x-1|K')}{f(x-1|K)f(x|K')} \\ &= \frac{(K-x+1)!\,(N-K-n+x-1)!\,(K'-x)!\,(N-K'-n+x)!}{(K-x)!\,(N-K-n+x)!\,(K'-x+1)!\,(N-K'-n+x-1)!} \\ &= \frac{K-x+1}{K'-x+1}\frac{N-K'-n+x}{N-K-n+x} \\ &> 1, \end{aligned}$$

assuming all terms appearing in this expression are positive.

For the MLRP in n, we have

$$\begin{aligned}
&\frac{f(x|n)f(x-1|n')}{f(x-1|n)f(x|n')} \\
&= \frac{(n-x+1)!\,(N-K-n+x-1)!\,(n'-x)!\,(N-K-n'+x)!}{(n-x)!\,(N-K-n+x)!\,(n'-x+1)!\,(N-K-n'+x-1)!} \\
&= \frac{n-x+1}{n'-x+1}\frac{N-K-n'+x}{N-K-n+x} \\
&> 1.
\end{aligned}$$

For the MLRP in N, we have a simpler expression:

$$\begin{aligned}
&\frac{f(x|N)f(x-1|N')}{f(x-1|N)f(x|N')} \\
&= \frac{(N-K-n+x-1)!\,(N'-K-n+x)!}{(N-K-n+x)!\,(N'-K-n+x-1)!} \\
&= \frac{N'-K-n+x}{N-K-n+x} \\
&< 1.
\end{aligned}$$

This completes the proof. But for much of the analysis we do in the paper, we are actually interested in a different relationship between the parameters. Before we showed that seeing a higher outcome is more indicative of a higher value of K. But if we see the same outcome, a lower value of n is more indicative of a higher value of K.

We will use the following MLRP of the hypergeometric distribution.

Lemma 1.6. For all $0 \leq K' < K \leq N$, the extended-real-valued likelihood ratio

$$\frac{p(x;M,K,N)}{p(x;M,K',N)} = \frac{\binom{K}{x}\binom{N-K}{M-x}}{\binom{K'}{x}\binom{N-K'}{M-x}}$$

is increasing in x over the range $[(K'+M-N)_+, M\wedge K]$, where either the numerator or denominator is positive.

Proof. Within this range, the likelihood ratio is zero if $x < (K+M-N)_+$ and infinite if $x > M \wedge K'$, so we restrict attention to the range

$$[(K+M-N)_+, M\wedge K'].$$

For integers x such that $(K+M-N)_+ + 1 \le M \wedge K'$, some algebra gives

$$\frac{p(x;M,K,N)}{p(x;M,K',N)}\Big/\frac{p(x-1;M,K,N)}{p(x-1;M,K',N)} = \frac{K-x+1}{K'-x+1}\cdot\frac{N-K'-M+x}{N-K-M+x} > 1,$$

so the proof is complete. □

1.7.4 Other Equilibria

Receiver-Optimal Equilibrium

Since payoffs depend only on the state and whether the product is purchased. Therefore, the payoff-relevant outcome is the buying function $b\colon \Theta \to [0,1]$, where $b(\theta)$ is the probability of buying given that the realized state is θ. By symmetry, payoffs depend only on $|\theta|$, not θ. Therefore, we can focus on the *buying vector* $b = (b_0, \dots, b_N)$, where b_j denotes the probability of buying conditional on $|\theta| = j$, which can be computed from the buying function:

$$b_j = \binom{N}{j}^{-1} \sum_{\theta:|\theta|=j} b(\theta). \tag{1.4}$$

Let $\tilde{v} : \{0, \dots, N\} \to \mathbb{R}$ such that $\tilde{v}(i) = v(\theta)$ if $|\theta| = i$. The payoffs from a buying vector b, under the state distribution π, are

$$u_S(b,\pi) = \sum_{j=0}^{N} \pi_j b_j, \qquad u_R(b,\pi) = \sum_{j=0}^{N} \pi_j b_j (\tilde{v}(j) - P).$$

Ideally we would like to characterize which buying vectors b can be induced by some equilibrium. A full characterization is not possible with our current methods, but we can obtain bounds on the equilibrium buying vectors by imposing the incentive constraints for the sender and the receiver.

Recall that the receiver can always choose an attribute to check randomly and can then buy if this randomly chosen attribute is good. This results in the buying vector $\bar{b} = (0, 1/N, 2/N, \dots, 1)$. Let $\bar{u} = u_R(\bar{b}, \pi)$. The receiver can also choose to never buy, yielding a utility of 0. A fundamental constraint on any equilibrium buying vector is that the receiver's utility must be at least $\bar{u}_+$. Of course this is a fairly crude lower bound. The sender's messages in equilibrium may reveal valuable information to the receiver that make new, more profitable deviations available.

The sender's incentive constraints prove much more useful. We can establish the following bound by considering a particular class of symmetric deviations by the sender. Namely, any type can uniformly mimic the types that have the same number of good attributes, plus one more.

Lemma 1.7. If a buying vector b in $[0,1]^{N+1}$ is induced by an equilibrium, then $b_{j-1} \geq (j-1)b_j/j$ for each $j \in [N]$.

Proof. Fix an equilibrium f, and let b denote the associated buying function. The buying vector is given in (1.4). Fix $j \in [N]$, and consider some state θ with $|\theta| = j$.

We can think of the state θ in two equivalent ways. As a vector, we have $\theta = (\theta_1, \dots, \theta_N)$, where θ_i equals 1 if the i-th attribute is good and 0 if the i-th attribute is bad. Alternatively, we can think of θ as a subset of $[N]$, where $i \in \theta$ if and only if the i-th attribute is good. In the following argument, we take the latter interpretation.

In the main text, we consider an abstract message space and define strategies in the standard way. That is, the sender chooses a map from states to messages, and the receiver chooses a map from messages to actions. This is a useful way to describe particular equilibria. For the following argument, it is helpful to apply some insights from the revelation principle of mechanism design.

Consider an equilibrium. By identifying each equilibrium message with the action it induces the receiver to take, we can represent an equilibrium by a stochastic mapping from states to actions. Intuitively, consider any arbitrary equilibrium. Then we can simply pool all the messages that induce the same action for the receiver, and then relabel each message by its induced action. Therefore, if (c, b^0, b^1) is an action profile of the receiver that is chosen with positive density in equilibrium, then we can regard (c, b^0, b^1) as a message sent by the sender in equilibrium as well.

By the sender's incentive constraint,

$$b(\theta) = \sum_{i \in \theta} c_i b_i^1 + \sum_{i \in [N] \setminus \theta} c_i b_i^0.$$

For each $\ell \in \theta$, the sender's incentive compatibility constraint at $\theta \setminus \{\ell\}$ can send this message (c, b^0, b^1) and the receiver will then buy with probability

$$\sum_{i \in \theta \setminus \{\ell\}} c_i b_i^1 + \sum_{i \in [N] \setminus (\theta \setminus \{\ell\})} c_i b_i^0 = b(\theta) - c_\ell (b_\ell^1 - b_\ell^0).$$

The sender's incentive compatibility constraint at $\theta \setminus \{\ell\}$ implies that

$$b(\theta \setminus \{\ell\}) \geq b(\theta) - c_\ell(b_\ell^1 - b_\ell^0).$$

Summing over $\ell \in \theta$ gives

$$\sum_{\ell \in \theta} b(\theta \setminus \{\ell\}) \geq jb(\theta) - \sum_{\ell \in \theta} c_\ell(b_\ell^1 - b_\ell^0) \geq (j-1)b(\theta),$$

where last inequality follows from comparing the sum here with the expression for $b(\theta)$ above, and using the nonnegativity of b^0 and b^1. Recall that θ was an arbitrary state with j good attributes. Summing over all such states gives

$$\sum_{\theta:|\theta|=j} \sum_{\ell \in \theta} b(\theta \setminus \{\ell\}) \geq (j-1) \sum_{\theta:|\theta|=j} b(\theta).$$

The sum on the left side includes every type with $j-1$ good attributes exactly $N-j+1$ times, so we have

$$(N-j+1) \sum_{\theta':|\theta'|=j-1} b(\theta') \geq (j-1) \sum_{\theta:|\theta|=j} b(\theta).$$

That is,

$$(N-j+1)\binom{N}{j-1} b_{j-1} \geq (j-1)\binom{N}{j} b_j.$$

Expanding the binomial coefficients and simplifying yields the desired inequality. □

Combining this lemma, which follows from the sender's incentive constraints, with the crude bound derived above from the receiver's constraint, we obtain the following theorem.

Theorem 1.3. *If a buying vector b in $[0,1]^{N+1}$ is induced by some equilibrium, then*

$$b_{j-1} \geq \frac{j-1}{j} b_j, \quad j = 1, \ldots, N,$$
$$\sum_{j=0}^{N} \pi_j b_j(\tilde{v}(j) - P) \geq 0 \vee \sum_{j=0}^{N} \pi_j (j/N)(\tilde{v}(j) - P).$$

Recall again that the first set of inequalities comes from the sender's ability to imitate higher types and the second inequality holds because the receiver can always deviate to the strategy of uniform checking.

Also note that the first set of inequalities is independent of the distribution π, the valuation function v, and the price P. Let B denote the set of vectors $b \in [0,1]^{N+1}$ such that

$$b_{j-1} \geq \frac{j-1}{j} b_j, \quad j = 1, \ldots, N.$$

Geometrically, the theorem states that any equilibrium buying vector must lie in B and additionally lie above some (π, v, P)-dependent hyperplane slicing through the set B.

By maximizing $\sum_{i=0}^{N} \pi_i b_i$ over the inequalities above, we get a linear program whose value is an upper bound on the sender's utility in any equilibrium. With a little trick, we can use this bound to show that the top equilibria are not Pareto dominated by any other equilibria.

Theorem 1.4. *Among all equilibria, the sender-optimal top equilibrium is strongly Pareto optimal.*

Proof. For each $j = 1, \ldots, N$, let

$$e_j = (0/j, 1/j, \ldots, (j-1)/j, j/j, 0, \ldots, 0).$$

Then each vector b in B can be expressed uniquely as

$$b = \sum_{j=1}^{N} \alpha_j e_j$$

for some coefficients $a_j \in [0,1]$ such that

$$\sum_{j'=j}^{N} (j/j')\alpha_j \leq 1.$$

Specifically, $\alpha_j = b_j - jb_{j+1}/(j+1)$ for each j.

With this parametrization, maximizing the sender's utility subject to a constraint on receiver's utility and vice versa is given by a collection, so this can be solved by a greedy algorithm. So we have an exact representation of the Pareto frontier, and this applies even to asymmetric equilibria as well.

Specifically, suppose $\bar{u} \geq 0$, so that the receiver is willing to buy with some positive probability under the uniform strategy. Let k_R be the smallest value of k for which $u_R(e_k, \pi) \geq \bar{u}$. It can be shown that the Pareto frontier of B is precisely given by mixtures of top-k and top-k' equilibria for consecutive k and k' between 0 and k_R, where top-0 is interpreted as always buying. All of these are not incentive compatible though. Let k_S be the smallest value of k such that

$$u_R(e_N + N^{-1}e_{N-1} + \cdots + (k+1)^{-1}e_k, \pi) \geq \bar{u}.$$

By construction $k_S \leq k_R$. The top equilibria are precisely the top-k equilibria for k between k_S and k_R. The top-k_R equilibrium is the receiver-optimal equilibrium (Indeed he could not do any better through commitment, by Glazer and Rubinstein (2004)). The top-k_S equilibrium is the sender's best top equilibrium, but all these equilibria are Pareto efficient among *all* the equilibria.

□

This theorem shows that the receiver-optimal equilibrium (optimal among all equilibria, symmetric or asymmetric) is a top equilibrium. However, we cannot rule out that the sender can do slightly better under some other equilibria, and indeed we can construct asymmetric equilibria where this is the case.

Asymmetric Equilibrium

In the top-k equilibrium, the sender is pointing to the k highest attributes out of all the N attributes. Now we let the sender point to the k highest attributes among a prescribed subset consisting of K attributes where $K \leq N$. We could allow the sender to randomly choose K attributes first but for intuition let's say the sender points to the k highest attributes out of the first K attributes and ignores the remaining $N - K$ attributes. Recall that there is a trade-off between the frequency of seeing a good signal and how strong that signal can be. By doing this, the sender makes a signal stronger when $K < N$ than $K = N$. We call this strategy top-k of K strategy and use (k, K) as a shorthand. Then (k, N) corresponds to the regular top-k strategy we characterized before.

To be clear about the intuition, let's see an example. Suppose $N = 4$ and consider two types of the sender $(0, 0, 1, 0)$ and $(0, 0, 1, 1)$ with $\Pr(\theta = (0, 0, 1, 0)) = p$ and $\Pr(\theta = (0, 0, 1, 1)) = q$. We claim the signal in $(1, N - 1)$ is stronger than $(1, N)$.

To see why, note that in $(1, N)$, sender $(0,0,1,0)$ will point to attribute 3 with probability 1 and sender $(0,0,1,1)$ will point to attribute 3 with probability 1/2. Denote the message "my highest attribute of all attributes is attribute 3" by m_1 and "my highest attribute of the first three attributes is attribute 3" by m_2. And hence, $\frac{\Pr(m_1|\theta=(0,0,1,0))}{\Pr(m_1|\theta=(0,0,1,1))} = \frac{1}{1/2} = 2$. However, under $(1, N-1)$, both types will point to the third attribute with probability 1. We have $\frac{\Pr(m_2|\theta=(0,0,1,0))}{\Pr(m_2|\theta=(0,0,1,1))} = \frac{1}{1} = 1$. Then the receiver's posterior is: in $(1, N)$, $\frac{\Pr(\theta=(0,0,1,0)|m_1)}{\Pr(\theta=(0,0,1,1)|m_1)} = \frac{\Pr(m_1|\theta=(0,0,1,0))}{\Pr(m_1|\theta=(0,0,1,1))} \times \frac{\Pr(\theta=(0,0,1,0))}{\Pr(\theta=(0,0,1,1))} = 2\frac{p}{q}$ and in $(1, N-1)$, $\frac{\Pr(\theta=(0,0,1,0)|m_2)}{\Pr(\theta=(0,0,1,1)|m_2)} = \frac{\Pr(m_2|\theta=(0,0,1,0))}{\Pr(m_2|\theta=(0,0,1,1))} \times \frac{\Pr(\theta=(0,0,1,0))}{\Pr(\theta=(0,0,1,1))} = \frac{p}{q}$. So the receiver is more biased to $(0,0,1,1)$ when pointed to the highest attribute out of the first three attributes. It turns out that for fixed k, the signal strength (k, K) is decreasing in K and for fixed K, the signal strength (k, K) is increasing in k, so we have a partial order.

With respect to equilibrium buying probability, top-k of K equilibrium is the intermediate case between top-k equilibrium and top-K equilibrium. If we go back to the numerical example we used in Section 1.3.3, we can clearly see the price threshold and equilibrium buying probability of top-1 of 2 equilibrium are exactly between top-1 equilibrium and top-2 equilibrium.

Recall that in that numerical example, we let $N = 3$, $n = 1$, $v(\theta) = |\theta|$ and $|\theta| \sim \text{Bino}(N, q)$ where $0 < q < 1$. In a top-1 of 2 equilibrium, the sender ignores one attribute and points to the highest attribute of the remaining two. The receiver just checks the recommended attribute and buys unless both of those two attributes are bad. Thus, the equilibrium buying probability is $1 - (1-q)^2$. And we know that the receiver's expected utility from equilibrium strategy should be weakly larger than that from checking a random attribute and buying if and only if that attribute is good. Therefore, the condition for the existence of top-1 of 2 equilibrium is $(1 - (1-q)^2)(q + \frac{2q}{1-(1-q)^2} - P) \geq q(2q + 1 - P)$. So the highest price to sustain a top-1 of 2 equilibrium is $\bar{P}_{1,2} = 1 + q$. Recall that the price threshold of top-1 and top-2 equilibrium is $\bar{P}_1 = \frac{2}{2-q}$ and $\bar{P}_2 = \frac{1+3q}{1+q}$, respectively. Clearly, $\bar{P}_1 < \bar{P}_{1,2} < \bar{P}_2$ for all $0 < q < 1$. We plot the equilibrium buying probability and price thresholds in Figure 1.5a and Figure 1.5b, respectively. It is clear that top-1 of 2 equilibrium is an intermediate case between top-1 and top-2 equilibrium.

Hence, at $\bar{P}_{1,2}$ top-1 equilibrium is not sustainable. The sender's most-preferred top equilibrium is top-2 equilibrium. However, the sender is strictly better off in top-1 of 2 equilibrium which is an asymmetric equilibrium.

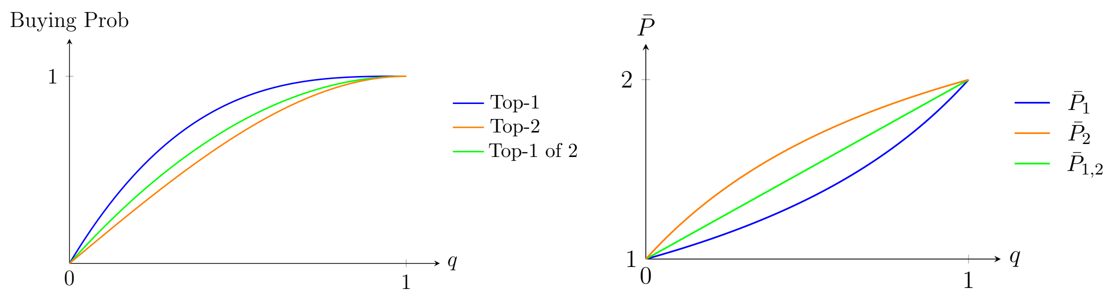

(a) Equilibrium buying probability **(b)** Price thresholds

Figure 1.5. Equilibrium Buying Probability and Price Thresholds

It is worth noting that the price threshold $\bar{P}_{1,2}$ in top-1 of 2 equilibrium coincides with the price $\hat{P}$ we used before (see Section 1.3.3). This implies we can replicate the outcome of top-1 of 2 equilibrium by a Bayesian persuasion with verification (BPV) setting. The intuition is as follows. It can be shown that at $\bar{P}_{1,2}$ (or $\hat{P}$), under BPV, the sender points to the highest attribute with probability 2/3 and the second highest attribute with probability 1/3. Under cheap talk, the sender will point to the highest attribute of two preselected attributes. It turns out that they are the same. Why? Suppose exactly one attribute is good (otherwise, the highest and second highest attributes are indistinguishable (either both one or both zero)). Under BPV, the sender will point to the good attribute with probability 2/3. Under cheap talk, the sender will point to the good attribute exactly when it is one of the two preselected attributes, which occurs with probability 2/3. Therefore, the buying probability as a function of $|\theta|$ is the same under both equilibria.

Other Symmetric Equilibria

This section constructs a parametric family of symmetric equilibria that are not top equilibria. We start with an example. Notice that it is sensitive to the parameters being "just right" in order to maintain indifference, so it is more fragile than the top equilibria. Moreover, in this example, there is a top equilibrium that induces a higher buying probability.

Suppose $N = 4$ and $n = 1$. The attributes are i.i.d. with success probability 1/2.

So that the probabilities of $|\theta|= 0, 1, 2, 3, 4$ are

$$1/16, 1/4, 3/8, 1/4, 1/16$$

respectively. Set $P = 7/4$. Consider the equilibrium where the receiver asks the sender to indicate the two highest attributes (unordered) and also the third-highest attribute. The receiver will check one of the three highest attributes and buy if and only if it is one, but he will check each of the two highest attributes with probability $c \in (1/3, 1/2)$ and he will check the third attribute with probability $1 - 2c$.

Next, we verify that this is an equilibrium. When the receiver checks one of the two highest attributes, his utility is

$$(1/4)(1/2)(1 - 7/4) + (3/8)(2 - 7/4) + (1/4)(3 - 7/4) + (1/16)(4 - 7/4) = 0.453$$

When the receiver checks the third highest attribute, his expected utility is

$$(1/4)(3 - 7/9) + (1/16)(4 - 7/4) = 0.453.$$

If the receiver picks randomly, then

$$(1/4)(1/4)(1 - 7/4) + (3/8)(1/2)(2 - 7/4) + (1/4)(3/4)(3 - 7/4) + (1/16)(4 - 7/4) = 0.375.$$

Therefore, this constitutes an equilibrium.

Ultimately, however, this is not of much interest. We know the sender does strictly better in a top-2 equilibrium. However, the top-1 equilibrium

$$2 - (15/16)(7/4) = 0.359.$$

is not available.

Now we construct this parametric family of symmetric equilibria.

Equilibria that only provide comparative information about the attributes can never increase the buying probability above a top equilibrium. Some absolute information is needed. For equilibria in this family, there are two types of messages. The first message indicates the top $\underline{k}_1$ attributes and then the next $\bar{k}_1 - \underline{k}_1$ attributes; the receiver chooses from among the highest attributes with probability c_1, and among the next $\bar{k}_1 - \underline{k}_1$ with complementary probability $1 - c_1$. The second message indicates the top $\underline{k}_2$ attributes and then the next $\bar{k}_2 - \underline{k}_2$ attributes; the receiver choose from among the highest attributes with probability c_2 and among the next $\bar{k}_2 - \underline{k}_2$ with probability $1 - c_2$. No matter the message and the attribute checked, the receiver buys if and only if he sees a one, except in the edge case where $\underline{k}_1 = 0$, in which case the receiver buys upon checking the "zeroth-highest" attribute.

Types $t \geq \bar{k}_2$ are indifferent about which message they send as buying is guaranteed either way. Among types $t < \bar{k}_2$, it turns out that "extreme" types (either below $\underline{t}$ or above $\bar{t}$) will choose to send the first message and "moderate" types (between $\underline{t}$ and $\bar{t}$ inclusive) will choose to send the second message.

The cutoff values $\underline{k}_1, \bar{k}_1$ for the first message, $\underline{k}_2, \bar{k}_2$ for the second message, and $\underline{t}$ and $\bar{t}$ for the types must satisfy the following inequality:

$$0 \leq \underline{k}_1 < \underline{t} < \underline{k}_2 < \bar{t} < \bar{t} + 1 < \bar{k}_1 < \bar{k}_2 \leq N.$$

This implies $N \geq 6$, and in order to make the first and last inequalities strict, we must in fact have $N \geq 8$. The inequalities can be explained as follows. We analyze the inequalities from the outside in. The outermost (weak) inequalities are trivial. The next inequalities $\underline{k}_1 < \underline{t}$ and $\bar{k}_1 < \bar{k}_2$ are needed so that the agents do not just split between the messages as high types and low types. then the inequalities $\underline{t} < \underline{k}_2 < \bar{t}$ ensure that the quality of the attribute the receiver checks after receiving the second message is uncertain. Finally, the inequalities $\bar{t} + 1 < \bar{k}_1$ ensures that seeing a zero among the lower attributes of the first message does not guarantee that the type is below $\underline{t} - 1$.

The continuous parameters are the price P and the interior probabilities

$$c_1, c_2, \pi_0, \ldots, \pi_N \in (0, 1),$$

such that

$$\sum_t \pi_t = 1.$$

Technically, c_1 and c_2 are equilibrium parameters, while P and (π_t) are parameters of the environment. We will see that the incentive constraints imply $c_1 < c_2$ and

$\underline{t} < P < \bar{t}$, but these constraints need not be included explicitly.

To simplify notation, let $\mathcal{M} = [\underline{t}, \bar{t}]$ and $\mathcal{E} = [0, \underline{t} - 1] \cup [\bar{t} + 1, \bar{k}_2 - 1]$, so $\mathcal{M} \cup \mathcal{E} = [0, \bar{k}_2 - 1]$. The types $t \geq \bar{k}_2$ will always buy in any equilibrium. Now we impose the equilibrium constraints.

- Attribute ordering:

$$c_1/\underline{k}_1 \geq (1 - c_1)/(\bar{k}_1 - \underline{k}_1), \qquad c_2/\underline{k}_2 \geq (1 - c_2)/(\bar{k}_2 - \underline{k}_2).$$

- Extreme sender:

$$c_1 + (1 - c_1)(\underline{t} - 1 - \underline{k}_1)/(\bar{k}_1 - \underline{k}_1) \geq c_2(\underline{t} - 1)/\underline{k}_2,$$
$$c_1 + (1 - c_1)(\bar{t} + 1 - \underline{k}_1)/(\bar{k}_1 - \underline{k}_1) \geq c_2 + (1 - c_2)(\bar{t} + 1 - \underline{k}_2)/(\bar{k}_2 - \underline{k}_2).$$

- Moderate sender:

$$c_2\underline{t}/\underline{k}_2 \geq c_1 + (1 - c_1)(\underline{t} - \underline{k}_1)/(\bar{k}_1 - \underline{k}_1),$$
$$c_2 + (1 - c_2)(\bar{t} - \underline{k}_2)/(\bar{k}_2 - \underline{k}_2) \geq c_1 + (1 - c_1)(\bar{t} - \underline{k}_1)/(\bar{k}_1 - \underline{k}_1).$$

- Receiver:

$$\sum_{t \in \mathcal{E}} (\tilde{v}(t) - P)\pi_t \left[\frac{t}{\underline{k}_1} \wedge 1 - \frac{(t - \underline{k}_1)_+}{\bar{k}_1 - \underline{k}_1} \wedge 1 \right] = 0,$$
$$\sum_{t \in \mathcal{M}} (\tilde{v}(t) - P)\pi_t \left[\frac{t}{\underline{k}_2} \wedge 1 - \frac{t - \underline{k}_2}{\bar{k}_2 - \underline{k}_2} \wedge 1 \right] = 0,$$

where $\tilde{v}(t) = v(\theta)$ such that $|\theta| = t$. A few changes are needed in the special case $\underline{k}_1 = 0$: the initial attribute ordering constraint is trivially satisfied (which is immediate upon expanding the product), and in the receiver's first constraint, $0/0$ must be replaced with 1.

Of course we also want to impose the constraint that the receiver does not want to always buy:

$$\sum_t (\tilde{v}(t) - P)\pi_t t/N > \sum_t (\tilde{v}(t) - P)\pi_t.$$

And we will want to compare to the top-k equilibria for $\underline{k}_1 < k \leq \underline{k}_2$. We will then impose the inequality that receiver cannot achieve the top-$(k-1)$ equilibria:

$$\sum_t (\tilde{v}(t) - P)\pi_t \frac{t}{N} > \sum_t (\tilde{v}(t) - P)\pi_t \frac{t}{k-1} \wedge 1.$$

Finally we will compare the buying probability in the top-k equilibrium

$$\sum_t \pi_t \frac{t}{k} \wedge 1$$

to the probability under the proposed equilibrium:

$$c_1 \sum_{t\in\mathcal{E}} \pi_t \frac{t \wedge \underline{k}_1}{\underline{k}_1} + (1-c_1) \sum_{t\in\mathcal{E}} \frac{(t-\underline{k}_1)_+}{\bar{k}_1 - \underline{k}_1} \wedge 1 \\ + c_2 \sum_{t\in\mathcal{M}} \pi_t \frac{t \wedge \underline{k}_2}{\bar{k}_2 - \underline{k}_2} + (1-c_2) \sum_{t\in\mathcal{M}} \pi_t \frac{t-\underline{k}_2}{\bar{k}_2 - \underline{k}_2} \wedge 1 + \sum_{t\geq \bar{k}_2} \pi_t.$$

Formally, we will maximize this difference subject to the constraint that the top-$(k-1)$ strategy profile is not an equilibrium (which implies that always buying is not an equilibrium).

Suppose $k \leq \underline{k}_2$. The gain over the top-k equilibrium is

$$c_1 \sum_{t\in\mathcal{E}} \pi_t \left[\frac{t\wedge \underline{k}_1}{\underline{k}_1} - \frac{t \wedge k}{k}\right] + (1-c_1)\sum_{t\in\mathcal{E}} \pi_t \left[\frac{(t-\underline{k}_1)_+}{\bar{k}_1 - \underline{k}_1} \wedge 1 - \frac{t\wedge k}{k}\right] \\ + c_2 \sum_{t\in\mathcal{M}} \pi_t \left[\frac{t \wedge \underline{k}_2}{\bar{k}_2 - \underline{k}_2} - \frac{t\wedge k}{k}\right] + (1-c_2)\sum_{t\in\mathcal{M}} \pi_t \left[\frac{t-\underline{k}_2}{\bar{k}_2 - \underline{k}_2} \wedge 1 - \frac{t \wedge k}{k}\right].$$

The key observation is that this reduces to a linear programming problem once P is fixed. Therefore, we will iterate over values of P, and for each fixed P solve the maximization problem. To make the inequalities weak, we will replace the strict inequalities with weak inequalities with a small tolerance $\varepsilon > 0$.

For fixed P, there are $N+2$ parameters. We will group the constraints according to equalities with zero on the right side, inequalities with zero on the right side, inequalities with ε and the RHS and one inequality with 1 on the RHS.

1.7.5 Correlation Between the Attributes

In this section, we are interested in how the pairwise correlation of the attributes affects the price threshold. Consider a conditional independence model for exchangeable binary random variables. Since each attribute follows the same Bernoulli distribution, suppose the common Bernoulli parameter p is random with cumulative distribution function F whose support is $[0, 1]$. Given p, the binary random variables $\theta_1, \ldots, \theta_N$ are conditionally i.i.d. We have

$$\Pr(\theta = \hat{\theta}) = \int_0^1 p^s(1-p)^{N-s} dF(p)$$

where $s = |\hat{\theta}| \in \{0, \ldots, N\}$.

Furthermore, assume F is a Beta(α, β) distribution, with density

$$f(p) = [B(\alpha, \beta)]^{-1} p^{\alpha-1}(1-p)^{\beta-1}$$

where $0 < p < 1$, $B(\alpha, \beta)$ is the beta function ($\alpha > 0, \beta > 0$) to ensure that the total probability integrates to 1. Then we have

$$\Pr(\theta = \hat{\theta}) = \frac{B(\alpha+s, \beta+N-s)}{B(\alpha,\beta)}.$$

where $s = |\hat{\theta}| \in \{0, \ldots, N\}$.

The pairwise correlation coefficient is

$$\rho = \frac{1}{\alpha+\beta+1}$$

which is a function of α and β. Therefore, we can vary the correlation of the attributes by changing the value of the distribution parameters α and β.

If we assume p follows a Beta(α, β) distribution, then $|\theta|$ has a Beta-binomial(α, β) distribution. The probability mass function is

$$\Pr(|\theta| = s) = \binom{N}{s} \frac{B(\alpha+s, \beta+N-s)}{B(\alpha,\beta)}$$

where $s = 0, \ldots, N$, and $\alpha, \beta > 0$. In particular, we have

$$\pi_0 = \Pr(\theta = \mathbf{0}) = \frac{B(\alpha, \beta+N)}{B(\alpha,\beta)}.$$

Moreover, by standard computation, we have

$$\mathbb{E}(|\theta|) = \frac{N\alpha}{\alpha+\beta} \quad \text{and} \quad \mathrm{Var}(|\theta|) = \frac{N\alpha\beta(\alpha+\beta+N)}{(\alpha+\beta)^2(\alpha+\beta+1)}.$$

Consider the special case in which $v(\theta) = |\theta|$. By Proposition 1.1, the price threshold reduces to

$$\bar{P}_1 = \frac{N\mathbb{E}(|\theta|) - \mathbb{E}(|\theta|^2)}{N(1-\pi_0) - \mathbb{E}(|\theta|)}.$$

Note that $\mathbb{E}(|\theta|^2) = (\mathbb{E}(|\theta|))^2 + \mathrm{Var}(|\theta|)$. Hence, we can rewrite $\bar{P}_1$ as a function of only α and β. However, we are not able to have a closed form of the price threshold as a function of the correlation coefficient ρ because of the Beta functions. We will plot the relationship between $\bar{P}_1$ and α and β numerically.

In the following numerical example, let $N = 10$. Since the values of the two parameters α and β control the shape of the Beta distribution, we start with the case that the Beta density function is symmetric about 1/2, i.e., $\alpha = \beta$. In particular, we vary their value between 1 and 3 which corresponds to symmetric unimodal Beta density functions. Notice that the pairwise correlation coefficient ρ is decreasing in α and β. But as Figure 1.6a shows, as α increases, the price threshold is increasing when the attributes become less and less correlated. Next, we fix β at 2 and increase α from 1 to 3. The Beta density function now is left skewed, i.e., a negatively skewed distribution. Figure 1.6b shows similar effect of the increase of α on the price threshold as the first case. Third, we fix α at 2 and increase β from 1 to 3. But this is just the mirror image (the reverse) of the Beta density function curve in the second case. And this is clear as Figure 1.6c illustrates. At last, we allow for simultaneous change of α and β. The change of the price threshold in this case is ambiguous since the effects of the increase of α and β offset. See Figure 1.6d. Of course, ρ also decreases when both α and β go up.

We see that when β is fixed, increasing α increases the price threshold, as the correlation decreases. When α is fixed, increasing β decreases the price threshold, as the correlation decreases. It seems a bit counter-intuitive that when the attributes are less correlated, the receiver is more willing to pay a higher price. However, notice that the difference between the effects of α and β arises from the (ex ante) expected quality $\mathbb{E}(|\theta|)$. Increasing α increases the (ex ante) expected quality while increasing β decreases the (ex ante) expected quality. Different pairs of α and β can even result in the same correlation coefficient but different values of the expected quality. So in effect many things are changing at the same time. Intuitively, higher expected quality and higher correlation each tend to increase the price threshold. But in our numerical example the mean-quality and correlation effects are offsetting each other.

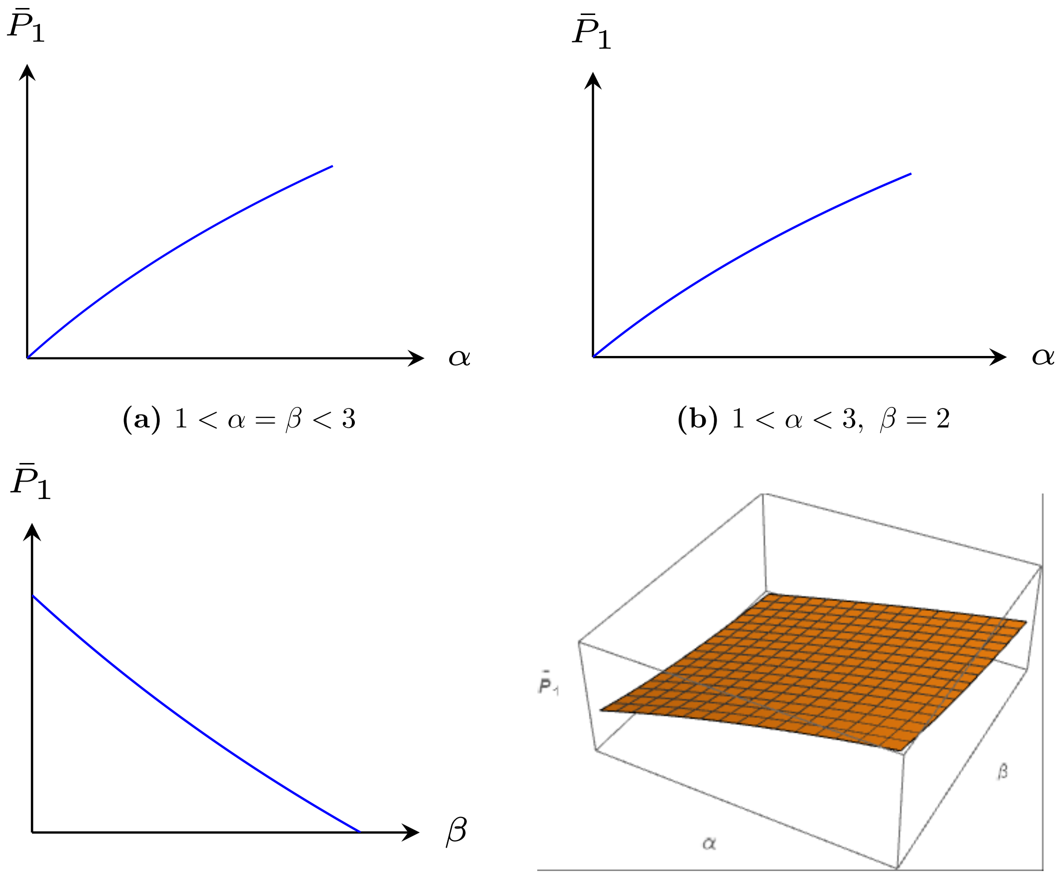

(a) $1 < \alpha = \beta < 3$ **(b)** $1 < \alpha < 3,\ \beta = 2$

(c) $\alpha = 2,\ 1 < \beta < 3$ **(d)** $1 < \alpha < 3,\ 1 < \beta < 3$

Figure 1.6. Price Threshold and Correlation

Chapter 2

Biased and Uninformed: Delegating to Encourage Information Acquisition

2.1 Introduction

Organizations are arranged into different branches that control different affairs. For example, it is impossible for a CEO to make every decision within a firm. Instead, she must delegate more routine decisions to lower-level managers. A crucial trade-off arises between control and information. The CEO would like to give more discretion to the subordinate managers in order to encourage more initiatives. However, at the same time she also has the fear that the managers may be empire builders who implement projects that are undesirable from the organization's point of view. At first sight, it seems easy to solve the incentive problem by simply providing contingent monetary transfers. But in practice, the manager's salary does not directly depend on some specific decisions he makes but rather depends on the whole performance of the organization. Instead of monetary incentives, the CEO just specifies what decisions the manager can or cannot make and allows him to take any action he wants within some admissible set.

This common trade-off illustrates the main features of our model. We consider an uninformed principal who delegates decision-making to an *initially* uninformed agent. There is some conflict of interest between them in the sense that the principal's ideal decision may differ from that of the agent. The agent can learn the payoff-relevant state by exerting costly effort. The principal cannot seek out information herself, due to, say, lack of time.

The timing of the game is as follows. First, the principal commits to an admissible set within which the agent is free to choose any decision. Second, the agent decides how much effort to exert to learn the state. At last, the agent takes an action from the set that the principal prescribed and payoffs are realized.

The principal cannot overrule the delegation set after she commits. There is no renegotiation between the two parties and outcome-contingent transfers are not allowed. By ruling out contingent transfers, the principal can influence the agent's information acquisition decision and action choice only by the prescribed admissible set.

The fundamental trade-off the principal faces here is between counteracting the agent's bias by limiting leeway for him and incentivizing the agent to acquire information by giving discretion to him. To illustrate this, on the one hand, standard delegation literature shows that more discretion should be given to a more aligned agent (a.k.a "Ally Principle"). On the other hand, it is intuitive to see more discretion tends to encourage effort input. For instance, a teacher who knows that he has to award the students either "pass" or "fail" has a smaller incentive to learn about the true performance of the students than a teacher who is allowed to award any grade, and an administrator who knows that he has to choose an either conservative or liberal policy has a smaller incentive to learn about what is the most appropriate decision on behalf of the society.

The aim of this paper is to characterize the optimal delegation set given the trade-off above. We use the canonical framework where the principal and the agent have quadratic loss utility function and the state follows uniform distribution over the unit interval. We find that if the cost of information acquisition is very small so that the agent will learn the state anyhow, then by the standard delegation literature, under uniform distribution, the optimal delegation set is an interval that shrinks as the agent's bias goes up.

If the information acquisition cost is not trivial, then the optimal delegation set is either a connected interval or a disconnected set that consists of two disjoint intervals. Specifically, the optimal delegation set is parameterized by two variables: the upper bound and the radius of the hole around the agent's most preferred uninformed decision. The former restricts the agent's expression of bias while the latter characterizes the punishment on the agent if he does not acquire information.

In order to explicitly solve for the optimal delegation set, we use the quadratic cost function to compute the optimal parameter values. We separate into cases according to the form of the optimal delegation set (connected or disconnected) and whether the agent acquires information fully. In each case, we analytically characterize the exact values of the delegation set.

Specifically, if the principal finds it optimal to induce full information acquisition and the optimal delegation set is a connected interval, then as the effort cost and the bias increase, the principal must enlarge the delegation set to motivate the agent to acquire information. In contrast, if full information acquisition is not optimal and the optimal delegation set is still connected, then the delegation set will shrink when the effort cost and the bias increase. For larger effort cost, the agent will acquire information with very low probability, so the principal would like to force the agent to take the principal's optimal uninformed decision by reducing the upper bound towards it. The reason that the upper bound is decreasing in the bias is clear. In this case, the bias is relatively high. The principal would rather give the agent less discretion as the bias increases.

If maximal effort input is optimal and the optimal delegation set is disconnected, then as the agent becomes more biased or information becomes more costly to acquire, the principal must distort the delegation set by more to incentivize information acquisition. Now the principal has two instruments available to encourage information acquisition: increasing the upper bound and hollowing out a larger interval around the agent-preferred uninformed decision.

Interestingly, the agent's bias can help the principal in the sense that the principal can punish the agent without punishing herself if the agent does not learn the state. Indeed, when the agent is perfectly aligned, introducing this hole will hurt the principal since her optimal no-information decision is the same as the agent. This makes the intuition from the classical delegation problem, which suggests that a principle would always prefer to delegate to a less biased agent, no longer hold. We find that the principal's payoff is highest when the agent's bias is nonzero.

Finally, if the optimal information acquisition is not full and the optimal delegation set is disconnected, then as effort cost goes to infinity, the agent will learn the state with very low probability and the principal would like the agent to take the principal's best no-information decision by enlarging the hole towards it.

Related Literature

Holmström (1977, 1984) first developed the framework for the constrained delegation problem in which an uninformed principal specifies a set of decisions from which a perfectly informed but biased agent is allowed to freely take. He shows the existence of a solution to this problem under very general conditions. By assuming interval admissible set, he shows that the more aligned the agent is, the more discretion should be awarded. In the subsequent works, one strand of the literature has gone into finding the weakest conditions that guarantee the optimality of interval delegation. For example, Martimort and Semenov (2006), Alonso and Matouschek (2008), and Amador and Bagwell (2013).

Another strand of the literature focuses on characterizing the optimal delegation set. This is first done by Melumad and Shibano (1991) under the quadratic-uniform framework. They find the shape of the optimal delegation set crucially depends on the relative sensitivity of the players' ideal decision to the state: if the sensitivity is similar a single interval is optimal, otherwise there can be gaps in the optimal permission set. Alonso and Matouschek (2008) further confirm this insight with more general distributions and utility functions. Although the optimal delegation set in these papers may also feature a "hole," it plays a different role compared to our paper. In Melumad and Shibano (1991) and Alonso and Matouschek (2008), the "hole" is to induce more state-sensitive decision-making by an otherwise unresponsive agent. In our paper, the "hole" is used to incentivize information acquisition.

In the Holmström tradition models, the action choices of the agent can be contracted upon and the agent is assumed to be exogenously informed of the payoff-relevant state. Starting from Aghion and Tirole (1997), some other papers have investigated situations in which the principal faces an initially uninformed agent who first needs to exert effort to acquire information before making a decision. As the seminal paper that introduced the trade-off between loss of control and loss of initiative, Aghion and Tirole (1997) assume the actions of the agent cannot be contracted upon so that the principal either keeps control or fully delegates the decision-making to the agent. They show that in order to improve the agent's initiatives of information acquisition, full delegation can be optimal.

The closest paper to us is Szalay (2005), who also studies a delegation problem in which the agent is potentially informed in the sense that he can learn the state by incurring an effort cost, but the principal and the agent are perfectly aligned. Since there is no conflict of interest between the principal and the agent, the only concern of the principal is to incentivize the agent to acquire information. To this end, he shows that the principal can just simply force the agent to take extreme actions by ruling out the intermediate ones. In other words, the principal creates a gap around the prior optimal decision as a punishment if the agent does not acquire information. This is in stark contrast to the standard insight in Holmström (1977, 1984) where it is optimal to remove the extreme decisions for the principal who faces a biased agent with exogenous information about the state.

Our paper takes them as benchmarks and combines them together. That is, we consider a more general setting where the agent is biased and his information is endogenous. Naturally, in the limit, our results reduce to those they got. Nevertheless, we show that in general the principal has to restrict both the intermediate and extreme actions given the new trade-off of the principal. To solve the moral hazard problem, the principal removes a symmetric interval around the agent's most preferred uninformed decision. To refrain the agent from taking extreme actions, the principal removes an interval on top.

In other contexts, it has been observed that bias can be leveraged to encourage information acquisition. Li (2001) and Gerardi and Yariv (2008) are closest to our paper in that they point out the potential for an undesirable default option to encourage information acquisition. Che and Kartik (2009) and Argenziano et al. (2016) both study endogenous information acquisition in communication games. There bias can be helpful because it motivates the sender to better acquire information in order to change the decision-maker's action.[1]

In our setting, the principal delegates authority to the agent, so the agent directly chooses the decision. We show that even in this context, bias can help the principal. Crucially, the principal chooses a *different* delegation set. For a fixed delegation set, increasing bias would always make the principal worse off, but by adjusting the delegation set, the principal can do better.

[1] In a legal context, Bubb and Warren (2014) study optimal regulation, but the delegation setting is different. In particular, they do not study the form of the delegation set. See also Malcomson (2009) and Demski and Sappington (1987) for a general contract-theory approach to the problem of encouraging information acquisition.

Since we follow the incomplete contracting approach (as in Grossman and Hart (1986) and Hart and Moore (1990)) in which contingent monetary transfers are ruled out, the only incentive device the principal can use is the discretion awarded to the agent. Another related paper by Armstrong and Vickers (2010) also features the interaction between the delegation sets and the responses that they induce in one variant of their main model. They study a stylized model in which there are a finite number of available projects and the payoffs of the principal and the agent from the project chosen by the agent are common knowledge. The agent only has private information about what those possible projects are. They show that when the principal can influence the likelihood of finding a project, the principal allows some projects that are detrimental to her own interests.

Instead of an *initially uninformed* agent, a recent paper by Semenov (2018) considers a delegation problem with a *possibly informed* agent, i.e., the agent knows the state with some exogenous probability. Without the moral hazard problem that we have, he shows how, depending on the preference alignment, a disconnected delegation set may be optimal for the principal. He also briefly discusses the endogenous information case but the focus is mainly on how the agent's incentives to acquire information change with the discretion he obtains.

The rest of this paper is organized as follows. Section 2.2 presents the model. Section 2.3 characterizes the form of the optimal delegation sets. Section 2.4 shows that the principal may prefer biased to unbiased agents. In Section 2.5, we explicitly solve for the optimal delegation set with quadratic costs, and then analyze comparative statics. Section 2.6 concludes. Proofs are in Section 2.7.

2.2 Model

2.2.1 Setting

There are two players: a principal and an agent. The principal controls a decision $y \in \mathbf{R}$. Payoffs from the decision depend a state θ, which is drawn from the uniform distribution on $[0, 1]$. The principal and the agent have conflicting quadratic loss preferences. Their utilities from decision y in state θ are

$$u_P(y, \theta) = -(y - \theta)^2, \qquad u_A(y, \theta) = -(y - \theta - \beta)^2.$$

Here β is the agent's *bias*. Without loss, we assume $\beta \geq 0$.

Initially, neither player observes the state realization. But the agent can privately experiment, at a cost, in order to learn the state. Following Szalay (2005), we assume that the agent chooses how much effort to put into experimentation. This effort level determines the probability that the experiment is a success. If the experiment succeeds, the state is perfectly revealed to the agent. Otherwise, the experiment fails, and the agent learns nothing about the state. The principal does not observe the agent's effort choice, the outcome of the experiment, or the state realization.

We normalize the effort level so that it coincides with the probability that the experiment is a success. Thus, the agent chooses effort $e \in [0, 1]$. The agent's cost from exerting effort e is given by a twice continuously differentiable cost function

$$c\colon [0, 1) \to \mathbf{R}$$

satisfying $c(0) = 0$, $c'(0) = 0$, and $c''(e) > 0$ for $e \in (0, 1)$. Set $c(1) = \lim_{e \uparrow 1} c(e) \in [0, \infty]$. If $c(1)$ is finite, then full information acquisition is possible.[2]

The principal delegates decision-making to the agent. This is equivalent to committing to a deterministic map from type reports to decisions. In particular, the principal cannot commit to contingent transfers or stochastic mechanisms. We assume the agent participates. These assumptions are standard in the delegation literature, and seem reasonable for many decisions made within organizations.[3]

The timing is shown in Figure 2.1. First, the principal selects a compact delegation set $D \subset \mathbf{R}$. The agent observes the delegation set D and chooses effort $e \in [0, 1]$. With probability e, the experiment is a success. In this case, the agent privately observes the state realization and then selects a decision from D. With probability $1 - e$, the experiment is a failure. In this case, the agent privately observes that the experiment has failed and then selects a decision from D. Finally, payoffs are realized.

The principal selects a compact delegation set to maximize her expected utility. Each choice of delegation set induces a two-stage optimization problem for the agent. We first analyze the agent's problem, and then turn to the principal's.

[2] In particular, we do not impose Szalay's (2005) condition that $c'(e) \uparrow \infty$ as $e \uparrow 1$. Therefore, we do not rule out the corner solution $e = 1$ in the agent's effort choice problem.

[3] For example, Alonso and Matouschek (2008) give a few justifications for these restrictions. Transfers may be forbidden by law; the principal may not be able to observe the agent's decision and payoff; or it may be more economical not to commit to a complicated transfer scheme or stochastic mechanism.

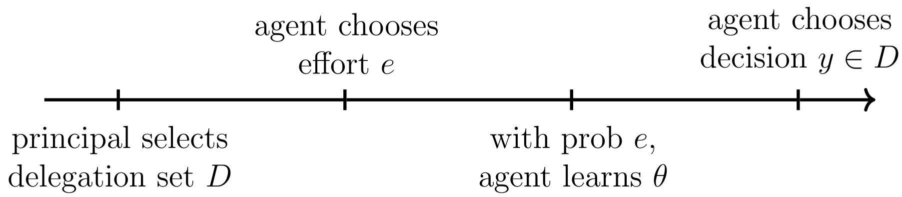

Figure 2.1. Timing

2.2.2 Agent's Problem

Suppose the principal has selected a compact delegation set D. The agent first chooses how much effort to exert and then selects a decision from D. We solve the agent's problem backwards, starting with his decision choice. If the experiment is a failure, then the agent solves

$$\max_{y \in D} \mathbf{E}[u_A(y, \theta)].$$

Here and below, all expectations are taken over the state θ. Since D is compact, this problem has a solution. Let $y_0^*(D)$ denote the agent's optimal decision, with the convention that ties are broken in the principal's favor.[4] We call $y_0^*(D)$ the agent's *uninformed decision* from delegation set D.

If the experiment is a success, and the agent observes the realized state θ, then he solves

$$\max_{y \in D} u_A(y, \theta).$$

Let $y_1^*(D, \theta)$ denote the agent's optimal decision, with ties broken in the principal's favor, as above. This *informed decision* depends on the state realization as well as the delegation set.

To set up the agent's effort choice problem, we introduce further notation. Denote the agent's expected utilities by

$$u_{A,0}(D) = \mathbf{E}[u_A(y_0^*(D), \theta)], \qquad u_{A,1}(D) = \mathbf{E}[u_A(y_1^*(D, \theta), \theta)].$$

The subscript number indicates whether the agent observes the state. The agent's

[4]The agent's expected utility is strictly quasiconcave in his decision, so there is an unique maximizer if D is an interval. But we will need to work with delegation sets that are not intervals. This is where tie-breaking has bite. The agent may be indifferent between two decisions, and the principal will have a strict preference as long as β is positive. If $\beta = 0$, choose $y_0^*(D)$ arbitrarily.

expected utility gain from observing the state is denoted by

$$\Delta_A(D) = u_{A,1}(D) - u_{A,0}(D).$$

Information can only help the agent, so $\Delta_A(D) \geq 0$ for every delegation set D.

Now the agent's effort choice problem is

$$\max_{e \in [0,1]} u_{A,0}(D) + e \cdot \Delta_A(D) - c(e). \tag{2.1}$$

Since c is strictly convex, this problem has an unique solution, which we denote by $e^*(D)$. The agent's marginal benefit from exerting effort is constant and equal to the utility gain $\Delta_A(D)$. The marginal cost $c'(e)$ is strictly increasing. If $e^*(D) < 1$, the first-order condition is

$$c'(e^*(D)) = \Delta_A(D).$$

2.2.3 Principal's Problem

Now we turn to the principal's problem. If the principal offers delegation set D, then from the solution of the agent's problem, we can compute the principal's expected utility conditional on each outcome of the experiment. Let

$$u_{P,0}(D) = \mathbf{E}[u_P(y_0^*(D), \theta)], \qquad u_{P,1}(D) = \mathbf{E}[u_P(y_1^*(D, \theta), \theta)].$$

Recall that $e^*(D)$ denotes the agent's effort choice, given delegation set D. Putting all this together, the principal's delegation problem is

$$\max_{D} e^*(D)u_{P,1}(D) + (1 - e^*(D))u_{P,0}(D), \tag{2.2}$$

where the maximization is over all compact subsets D of $\mathbf{R}$. The principal is maximizing over sets rather than points, but in a suitable topology, it can be shown that the principal is maximizing an upper semicontinuous function over a compact set, and hence a solution exists.

Lemma 2.1 (Existence)**.** The principal's problem (2.2) has a solution.

2.3 Characterizing the Optimal Delegation Set

2.3.1 Benchmark: Delegating to an Informed Agent

Before turning to the main analysis, we first consider the classical delegation problem with an informed agent. With our notation, that problem can be expressed as

$$\max_D u_{P,1}(D), \tag{2.3}$$

where the maximum is taken over all compact subsets D of $\mathbf{R}$. The principal's value from this problem provides a convenient upper bound.

Lemma 2.2 (Informed-agent upper bound)**.** The principal's value from delegating to an informed agent (2.3) is weakly greater than her value from delegating to an initially uninformed agent (2.2).

The proof is short. Under the solution to the informed problem, the principal must weakly benefit (in expectation) when the agent learns the state; otherwise the principal could offer a singleton delegation set including only the agent's uninformed best response. When the agent is informed, the principal gets this weakly higher informed payoff with probability one.

Recall that the optimal delegation set with an informed agent is[5]

$$D^* = \begin{cases} [\beta, 1-\beta] & \text{if } \beta \leq 1/2, \\ \{1/2\} & \text{if } \beta > 1/2. \end{cases}$$

If $\beta = 0$, then the interests of the principal and the agent are perfectly aligned, so the principal offers the agent complete discretion. As the bias β increases, the principal imposes a *cap* of $1-\beta$ on the agent's action in order to restrict the agent's bias. The agent will not choose decisions strictly smaller than β, so decisions will come from the interval $[\beta, 1-\beta]$.

As β increases, the cap tightens. When the bias β equals $1/2$, the agent is too conflicted for the principal to benefit at all from his private information, and the principal offers a delegation set that results in the constant decision $1/2$, no matter

[5] See Holmström (1977) section 2.3.2 or Alonso and Matouschek (2008) section 8.1. Technically, it is the induced decision rule is unique, but we are following the convention of associating to each decision rule the minimal delegation set that induces it.

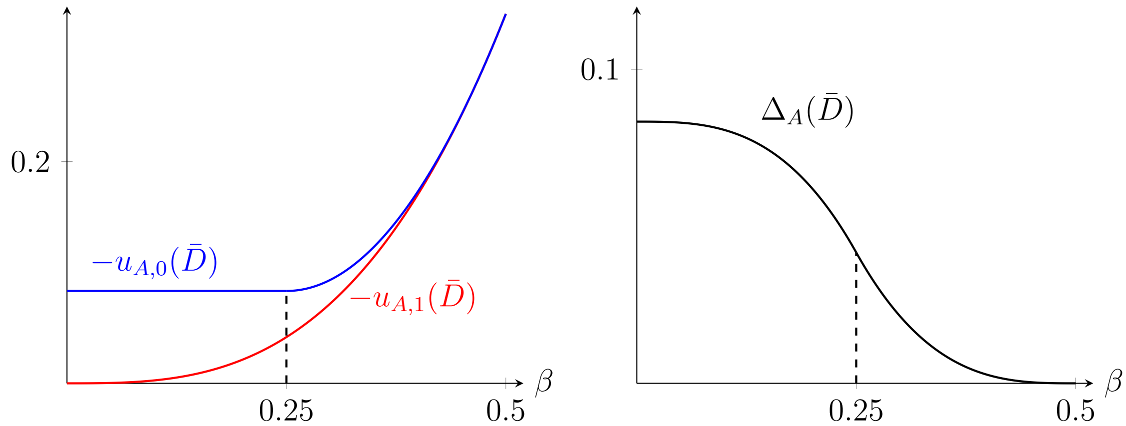

Figure 2.2. Agent's Payoffs from Delegation Set $\bar{D} = [\beta, 1-\beta]$

the state. With costly information acquisition, the principal certainly cannot benefit from the agent's private information, so the solutions will coincide. Formally, this follows from the upper bound in Lemma 2.2 since this bound can be achieved by setting $D = \{1/2\}$. To avoid this trivial case, we assume hereafter that $\beta < 1/2$.

Returning to our setting of an initially uninformed agent, suppose the principal selects this delegation set $\bar{D} = [\beta, 1-\beta]$.

Figure 2.2 shows the agent's expected utility gain $\Delta_A(\bar{D})$ from learning the state. The left panel plots the agent's expected losses with and without information. If the agent does not learn the state, then he matches his decision with the expectation $\mathbf{E}[\theta + \beta] = 1/2 + \beta$. For $\beta \leq 1/4$, this expectation lies inside the delegation set $[\beta, 1-\beta]$. In this case, the decision is not systematically biased from his bliss point, so his loss is simply $1/12$, the expected loss due to the variance of the state. Once the bias β crosses $1/4$, this expectation lies outside the delegation set $[\beta, 1-\beta]$, and hence the agent's loss increases in β. The agent's expected loss when he does learn the state is monotonically increasing in the bias, both because the cap $1-\beta$ shifts and also because the cap is effectively more restrictive as the agent's preferred action moves to the right.

The right panel of Figure 2.2 plots $\Delta_A(\bar{D})$, the difference in the agent's expected utility with and without information. The more aligned are the interests of the agent and the principal, the more the agent stands to gain by learning the state. If $c'(1) \leq \Delta_A(\bar{D})$, then the agent will choose $e = 1$ and thus learn the state perfectly. That is, though the agent is initially uninformed, he can acquire information so cheaply that the principal need not change the delegation set from the classical solution $[\beta, 1-\beta]$, so the principal gets the same payoff as if the agent had been initially informed. Therefore, the delegation set $\bar{D}$ achieves the upper bound in Lemma 2.2 and hence must be optimal.

It is natural to rule this case out by imposing an Inada condition on the cost function. We do not impose this restriction, however, because it would rule out quadratic cost, which we study below. While the solution is not new, this case tells us how costly information acquisition can become before it affects the delegation problem.

If $c'(1) > \Delta_A(\bar{D})$, then the principal cannot achieve the upper bound. That is, the cost of information acquisition reduces the principal's payoff. In particular, this is always the case if the marginal cost of effort is unbounded as effort e tends to 1. If the principal cannot achieve the informed-agent bound, then the principal is motivated to distort the delegation set away from $\bar{D}$ in order to encourage information acquisition. In the next section we characterize the form that this distortion takes.

2.3.2 Main Characterization

How can the principal encourage the agent to exert effort to acquire information? The agent's return to learning the state is captured by the utility difference $\Delta_A(D) = u_{A,1}(D) - u_{A,0}(D)$. Thus, there are two channels available. The principal can add new decisions to the delegation set in order to increase the agent's informed payoff. The principal can also reduce the agent's no-information utility by removing from the delegation set decisions near $1/2 + \beta$, the agent's most preferred choice without information.

Once this hole around $1/2 + \beta$ is removed, all other holes should be filled in by the standard intuition from the quadratic-uniform delegation. Interior holes do not shift the agent's average action but simply add further noise to the agent's decision. These observations give the following characterization.

An obstacle to uniqueness of the delegation set is that it is always possible to make changes to the delegation set that do not affect the decision rule. Therefore, we adopt the convention that all delegation sets are minimal. More explicitly, for each delegation set D, define the set of induced decisions by

$$Y^*(D) = \begin{cases} \{y_0^*(D)\} & \text{if } e^*(D) = 0, \\ \{y_0^*(D)\} \cup \{y_1^*(D,\theta) : \theta \in [0,1]\} & \text{if } 0 < e^*(D) < 1, \\ \{y_1^*(D,\theta) : \theta \in [0,1]\} & \text{if } e^*(D) = 1. \end{cases}$$

It follows that Y^* is idempotent, i.e., $Y^*(Y^*(D)) = Y^*(D)$ for all sets D. We adopt the convention of restricting to delegation sets D with the property that $Y^*(D) = D$. If $\beta < 1/2$, the optimal delegation set will always feature positive effort, so the second or third case will apply.

Our problem features both bias and information acquisition, which have previously been studied separately. In the standard delegation problem with bias, interval delegation is optimal. With information acquisition but not bias, it may be optimal to include a hole centered at 1/2 in order to encourage information acquisition. The characterization below shares both of these features. Below we will use the notation $B_r(x)$ to denote the open interval centered at x with radius r.

Theorem 2.1 (Characterization)**.** *Suppose* $\beta \in (0, 1/2)$ *and* $c'(1) > \Delta_A([\beta, 1-\beta])$. *Every optimal delegation set takes one of the following forms:*

1. $[\beta, \bar{y}]$ *for some cap* $\bar{y}$ *in* $(1/2, 1/2 + \beta]$*;*
2. $[\beta, \bar{y}] \setminus B_r(1/2 + \beta)$ *for some cap* $\bar{y}$ *in* $(\max\{1/2 + \beta, 1 - \beta\}, 1 + \beta)$ *and some radius* r *in* $(0, \bar{y} - 1/2 - \beta]$*;*
3. $[\beta, y_0] \cup \{h\}$ *for some default decision* y_0 *in* $(\beta, 1/2 + \beta)$ *and some high point* h *in* $(1 + y_0, 2 + 2\beta - y_0)$*.*

For the characterization, we assume that the cost of information acquisition is high enough to make a difference in the delegation problem. That is, if the principal offers the delegation set $[\beta, 1-\beta]$, the agent will not fully acquire information. With positive probability, the agent's experiment will fail, and the agent will select the decision in $[\beta, 1-\beta]$ that is closest to $1/2 + \beta$. Therefore, the principal's payoff is strictly smaller than the informed-agent upper bound. Now the principal must take into account how the delegation set affects the agent's incentive to experiment. In order to be optimal, a delegation set must take one of three forms,[6] termed *interval*, *hollow*, and *high-point* delegation, respectively.

[6]For some parameter values, there may be multiple delegation sets that are optimal. The theorem says that each of these sets must take one of three forms.

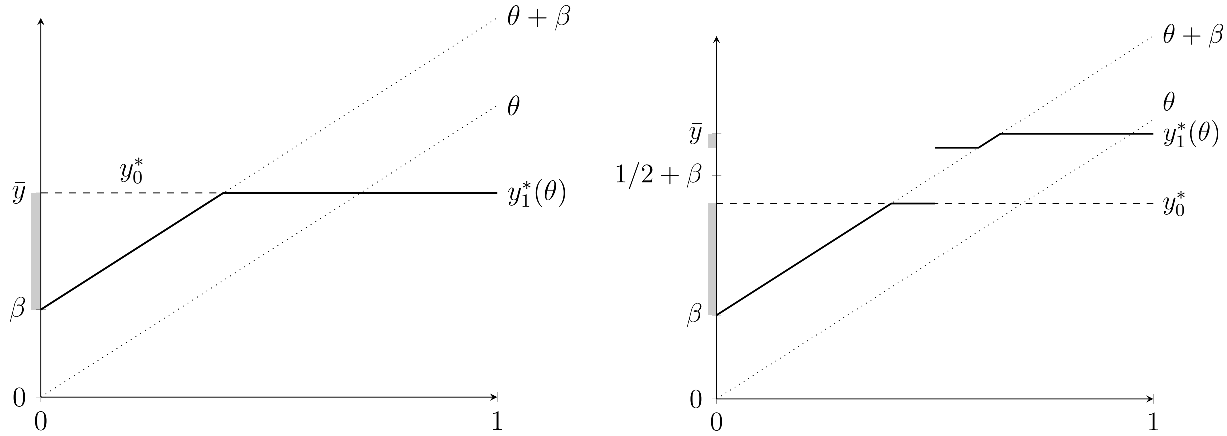

(a) Interval Delegation **(b)** Hollow Delegation

Figure 2.3. Optimal Delegation Set

Figure 2.3 illustrates interval and hollow delegation, and Figure 2.4 illustrates hight-point delegation. Note that y_0^* and $y_1^*(\theta)$ respectively denote the uninformed and informed decisions the agent will actually take given an optimal delegation set in each figure.

Interval delegation is familiar from the informed-agent delegation problem, but now the cap depends on the cost function as well as the agent's level of bias. Increasing the cap has two effects that are not present when the agent is informed. On the one hand, it increases the agent's utility gain from learning the state, so the agent exerts more effort. On the other hand, it shifts the agent's *uninformed* decision to the right. If the cap is below $1/2$, then both of these effects benefit the principal, so she does better by strictly increasing the cap. Thus the cap is above $1/2$. In this region, the two effects go in opposite directions, so $\bar{y}$ can be either larger or smaller than the cap $1 - \beta$ that is optimal when the agent is informed.

In case 2., the optimal delegation set is an interval with a hole around $1/2 + \beta$. Conditional on the agent learning the state, this hole makes both the principal and the agent worse off. But by eliminating the moderate decisions that are attractive to the agent if he is uninformed, this hole increases the agent's effort level and makes the principal better off. One particular possibility is that $1/2 + \beta + r = \bar{y}$. In this case, the optimal delegation set consists of an interval together with a single isolated point $\bar{y}$. This is not just as a knife-edge case, as we will show below in the case of quadratic costs.

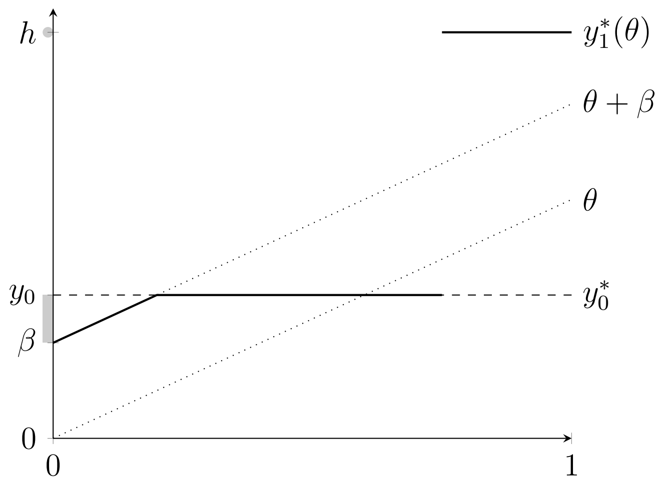

Figure 2.4. High-point Delegation

The final case 3. cannot be ruled out, but it requires quite extreme parameter values. Indeed, it takes some work to construct a cost function for which the optimal delegation set takes this form. An example is given in Section 2.7.7. In this case, the agent sometimes takes decisions that are dominated, not only from the principal's perspective, but from the agent's as well.

If the agent's bias β goes to zero, there is no conflict of interest between the principal and the agent. As shown in Szalay (2005, Proposition 1), the optimal delegation set takes the form

$$D^* = [0, 1] \setminus B_r(1/2),$$

for some $r \in [0, 1/2]$. While the *ex post* preferences of the principal and the agent are perfectly aligned, only the agent bears the cost of information acquisition. To encourage the agent to acquire information, the principal hollows out a *hole* around 1/2 to prevent the agent from taking safe decisions that perform reasonably well regardless of the state. When the agent is forced to take extreme actions, he is more motivated to acquire information.

Remark. Notice that we take advantage of the uniform distribution only when we show the principal is better off from closing an undesired hole or making an asymmetric hole around $1/2 + \beta$ symmetric conditional on the event that the agent learns the state, i.e., $u_{P,1}(D') > u_{P,1}(D)$. In fact, we can generalize the model to the case where the state follows a distribution such that $f(\theta) + \beta f'(\theta) > 0$ where $f(\theta)$ is the probability density function (clearly, uniform distribution over $[0, 1]$ satisfies this condition). To this end, let the cumulative distribution function of the state be $F(\theta)$ and the corresponding probability density function be $f(\theta)$ which is continuously differentiable and strictly positive for all $\theta \in [0, 1]$. Following Alonso and Matouschek (2008) define the *backward bias* as

$$T(\theta) = F(\theta)[(\theta + \beta) - \mathbf{E}[z|z \leq \theta]],$$

where z is a random variable distributed according to F. The backward bias characterizes the agent's weighted bias at a specific state θ conditional on the principal believing that the state is below θ.

Given this definition, by Lemma 7 in Alonso and Matouschek (2008), we can conclude that if the backward bias is strictly convex over the state space, then closing a hole or making a hole symmetric is better for the principal. Therefore, Theorem 2.1 still goes through.

So far we have characterized the *form* of the optimal delegation set, but we cannot directly compute the optimal set without choosing a particular functional form for the cost function. Section 2.5 considers the natural choice of quadratic cost, and solves for the optimal values of the parameters $\bar{y}$ and r. But first, we establish comparative statics that hold for any cost function satisfying our assumptions.

2.4 Comparative Statics in Bias and Information Cost

This section studies how the principal's value from delegation depends on the agent's bias and on the cost of information acquisition.

Theorem 2.2 (Comparative statics)**.**

1. *Fix the bias β and suppose the cost function c satisfies $c'(1) > \Delta_A([\beta, 1-\beta])$. If the marginal cost strictly decreases pointwise, then the principal's value from delegation strictly increases.*
2. *Suppose that for a fixed cost function c, it is optimal to induce an unbiased agent to acquire only partial information. Then the principal's payoff is maximized*

when the bias β is nonzero.

For the first part, we fix the bias and focus on the cost of information acquisition. The assumption rules out the case where information acquisition is so cheap that the principal's value is the same as if the agent were informed of the state. If information gets cheaper—in the sense that the marginal cost of information acquisition decreases pointwise—then the principal's payoff strictly increases. This conclusion follows from the fact that under optimal delegation, the principal must get a higher payoff when the agent is informed than when the agent is uninformed.

For the second part, we fix the cost function and consider changes in the bias β. Intuition from the classical delegation problem suggests that bias is bad for the principal, but it turns out that the level of bias that maximizes the principal's value from delegation is positive. The key to this result is that the principal chooses a different delegation set as the agent's bias increases. Starting at a bias level β_0 and using the delegation set D that is optimal at that bias level, the principal's payoff will always decrease as the bias increases. But, by adjusting the delegation set appropriately, the principal can increase her payoff.

The intuition is that including a hole around $1/2+\beta$ serves to punish the agent in the case that he does not observe the state. When the preferences of the principal and the agent are more aligned, a punishment for the agent is necessarily a punishment for the principal as well. But when the preferences of the principal and the agent diverge, the principal can better target her punishment on the agent without harming herself. Of course, at the moment the agent is making his decision, less bias is always better. Sometimes the first effect can dominate the second.

In our model, there is a single agent with a fixed level of bias. But imagine instead that there is a pool of agents with different levels of bias (but the same cost of information acquisition). According to our comparative static result, the principal's first choice will not be an unbiased agent. There is a strictly positive level of bias that the principal would strictly prefer.

2.5 Full Solution with Quadratic Costs

2.5.1 Quadratic Costs

In order to explicitly solve for the optimal delegation set, we must specify a cost function. We assume quadratic cost

$$c(e) = (1/2)\kappa e^2,$$

where κ parameterizes the cost of information acquisition. This parametrization satisfies the assumptions we made on the cost function, so the theorem applies. The marginal cost of effort is simply $c'(e) = \kappa e$. In particular, $c'(1) = \kappa$. With quadratic cost, the optimal effort level in the agent's problem (2.1) is

$$e^*(D) = \min\{1, \kappa^{-1}\Delta_A(D)\}.$$

That is, the agent's effort choice is increasing in the utility gain from learning the state, and decreasing in the cost coefficient κ, with the possibility of a corner solution $e = 1$.

Plugging the agent's optimal effort level into the principal's problem yields the principal's utility from delegation set D,

$$\begin{aligned}&\min\{1, \kappa^{-1}\Delta_A(D)\}u_{P,1}(D) + (1 - \min\{1, \kappa^{-1}\Delta_A(D)\})u_{P,0}(D) \\ &= u_{P,0}(D) + \min\{1, \kappa^{-1}\Delta_A(D)\}\Delta_P(D)\end{aligned}$$

where $\Delta_P(D) := u_{P,1}(D) - u_{P,0}(D)$.

To find the optimal delegation set, we separate into the cases from Theorem 2.1, optimize over the parameters within each regime, and then compare the solutions to find the global optimum. To determine which case we are in, we compute the optimal values separately for each case, and then compare. Therefore, the boundaries for different cases are given numerically.

Figure 2.5 plots the parameter space, with the bias β varying from 0 to 0.5 along the horizontal axis, and the cost κ of information acquisition varying from 0 to 0.25 along the vertical axis. The axes themselves correspond to the benchmark cases. The horizontal axis represents the classical delegation setting where cost vanishes. The vertical axis represents the case where the principal and the agent have perfectly

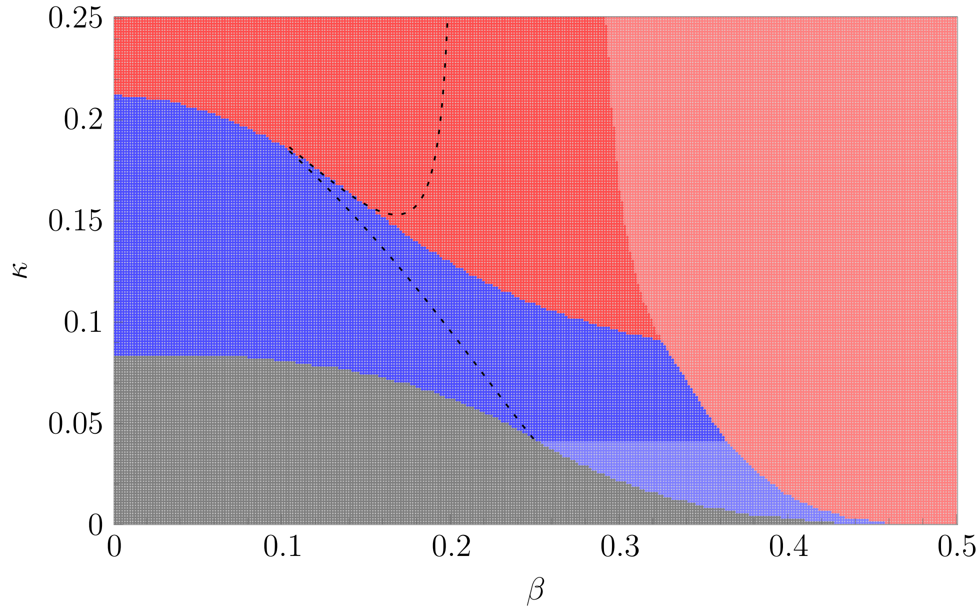

Figure 2.5. Parameter Space

aligned interests. The parameter space is partitioned into regions, which are discussed in detail below.

The parameter space is assigned three colors. In the gray region (low β and κ), the principal achieves the informed-agent upper bound. In the blue region (moderate β and κ), the principal distorts the delegation set away from the informed-agent delegation set in order to induce the agent to learn the state with probability one. This is called *full acquisition*. In the red region (high β and κ), the principal distorts the delegation set, but the agent still does not acquire information with probability one. This is called *partial acquisition*.

First we analyze the simplest case, the gray region, and then consider the other regimes in separate sections below. The gray region is defined by the condition $\kappa \leq \Delta_A(\beta)$, where $\kappa = c'(1)$ by construction, and $\Delta_A(\beta)$ denotes the agent's expected gain from information acquisition under the informed-agent optimal delegation set, as a function of the bias β. The graph shows that Δ_A is decreasing in β. In this region, the principal has no reason to distort the delegation set, and indeed the principal achieves the full payoff from the informed-agent upper bound.

2.5.2 Full Information Acquisition

Next we consider what happens as the cost κ increases, so that the agent will not fully acquire information under the informed-agent optimal delegation set. Then we are in the blue region, and the principal will distort the delegation set to encourage information acquisition. In this region, the optimal delegation set $D^* \neq [\beta, 1-\beta]$ but $e^*(D^*) = 1$.

The instruments available to the principal depend on the bias. If $\beta \leq 1/4$, then the agent's no-information bliss point $1/2+\beta$ is in $[\beta, 1-\beta]$. Therefore, the principal has two tools available. The principal can provide the agent with greater flexibility by increasing the cap $\bar{y}$ above $1-\beta$. The principal can also punish the agent for failing to acquire information by removing decisions in a hole centered around $1/2+\beta$. It turns out that the principal will use both instruments, which is clear from the first-order conditions. Increasing the radius has a third-order effect on the informed payoff but a second-order effect on the uninformed payoff.

Solving the first-order conditions, we find that in this case,

$$D^* = [\beta, 1 - \beta + 2r\beta] \setminus B_r(1/2 + \beta),$$

where r is positive and strictly increasing in both β and κ. In this case, the upper interval is non-degenerate and we are on the left hand side of the dashed line in the blue region.

If $\beta > 1/4$, then $1 - \beta < 1/2 + \beta$, so the agent's no-information bliss point is not available if given the delegation set $[\beta, 1 - \beta]$. In this case, the principal uses a single instrument, increasing the agent's flexibility, by pushing $\bar{y}$ above $1 - \beta$. Then we are in the light blue region. And we have

$$D^* = [\beta, \beta + (3\kappa)^{1/3}].$$

Note that the right endpoint is higher than $1-\beta$. Clearly, the length of this interval is strictly increasing in the cost κ of information acquisition. Plugging in $\kappa = (1/3)(1-2\beta)^3$ gives $1-\beta$. Thus, $\bar{y}$ increases continuously as κ increases, starting from $\bar{y} = 1-\beta$ at the boundary defined by $\kappa = (1/3)(1-2\beta)^3$. The upper limit $\bar{y}$ is strictly increasing in both β and κ. If the agent is more biased or information acquisition is more costly, the principal must distort the delegation set by more in order to encourage the agent to learn the state with probability 1. But why does the principal find it optimal to encourage maximal information acquisition? Indeed this is only true for κ small enough, but no matter the value of β, this will be true sufficiently near the boundary.

Finally, consider the case where the upper interval is degenerate. That is,

$$D^* = [\beta, 1/2 + \beta - r] \cup \{1/2 + \beta + r\},$$

where r depends only on κ and is strictly increasing in κ. This case corresponds to the dark blue region on the right hand side of the dashed line.

These results are formally summarized as follows.

Proposition 2.1 (Full information acquisition)**.** *Suppose $e^*(D^*) = 1$.*

1. ***Interval**: $D^* = [\beta, \beta + (3\kappa)^{1/3}]$.*
2. ***Hollow with upper point**: $D^* = [\beta, 1/2 + \beta - r] \cup \{1/2 + \beta + r\}$, where the radius r is independent of β and strictly increasing in κ.*
3. ***Hollow with upper interval**: $D^* = [\beta, 1 - \beta + 2r\beta] \setminus B_r(1/2 + \beta)$, where the radius r is strictly increasing in β and κ.*

2.5.3 Partial Information Acquisition

When β and κ are large enough, there will no longer be full information acquisition. Then we are in the red region. Again we can separate into cases according to whether the delegation set features one interval or two.

Interval In the light pink region, the optimal delegation set is a single interval $[\beta, \bar{y}]$. In this region, the cap $\bar{y}$ is strictly decreasing in both the bias β and the cost κ. Moreover, $\bar{y} < 1 - \beta$, so the agent chooses to bring the cap below the optimum level with an informed agent. This suggests that the benefits in the no-information case outweigh the costs.

Hollow with upper point In this case,

$$D^* = [\beta, 1/2 + \beta - r] \cup \{1/2 + \beta + r\},$$

where r is strictly increasing in κ. And we are on the right hand side of the dashed line in the dark red region. Many of the comparative statics in this region have ambiguous sign because the parameter r appears in many different places. To see this, consider the principal's payoff as a function of r,

$$U_P(r) = u_{P,0}(r) + \kappa^{-1}\Delta_A(r)\Delta_P(r).$$

Hollow with upper interval In this case,

$$D^* = [\beta, \bar{y}] \setminus B_r(1/2 + \beta) = [\beta, 1/2 + \beta - r] \cup [1/2 + \beta + r, \bar{y}],$$

where the second interval is non-degenerate, i.e., $\bar{y} > 1/2 + \beta + r$. This case corresponds to the red region on the left hand side of the dashed line. Now roles of r have been separated and we have hope for monotone comparative statics. Numerical computations suggest that $\bar{y}$ is strictly decreasing in β, as we might expect, but proving this analytically is challenging because there are two first-order conditions, both featuring high-degree polynomials.

The comparative statics in β on r are ambiguous. We can, however, give unambiguous comparative statics in κ: the radius r will move towards β as κ increases and converge to β as $\kappa \to \infty$.

To see the comparative statics in β on $\bar{y}$, consider the principal's payoff as a function of the cap $\bar{y}$ and the radius r:

$$U_P(\bar{y}, r) = u_{P,0}(r) + \kappa^{-1}\Delta_A(\bar{y}, r)\Delta_P(\bar{y}, r).$$

In this case, β is small enough such that $1/2 + \beta \in [\beta, 1 - \beta]$, so the principal has two instruments available to encourage information acquisition: increasing the upper bound $\bar{y}$ and hollowing out an interval of radius r about $1/2 + \beta$. Moreover, the first-order conditions require that

$$\bar{y} = 1 - \beta + 2r\beta,$$

so $\bar{y}$ and r must increase together, hence $\bar{y} > 1 - \beta$ and $r > 0$.

These results are formally summarized as follows.

Proposition 2.2 (Partial information acquisition)**.** *Suppose $e^*(D^*) < 1$.*

1. ***Interval**: $D^* = [\beta, \bar{y}]$, where the cap $\bar{y}$ is strictly below $1 - \beta$, is strictly decreasing in β and κ, and converges to $1/2$ as $\beta \to 1/2$ or $\kappa \to \infty$.*

2. ***Hollow***: *$D^* = [\beta, \bar{y}] \backslash B_r(1/2+\beta)$, where $|r-\beta|$ is decreasing in κ and converges to 0 as $\kappa \to \infty$.*

2.6 Conclusion

In many situations, it is too costly for the principal to make all the decisions. Hence, she may delegate some decisions to an agent. Instead of monetary incentives, the principal simply tells the agent what he can do and then awards the agent full discretion over this delegation set. We study how the principal uses this delegation set as the only incentive device to deal with two concerns: the agent's bias and moral hazard problem.

We have characterized the optimal delegation set under the canonical quadratic-uniform framework. It features a cap to limit distortion and a hole to encourage information acquisition. If the principal finds it optimal to induce full information acquisition, to encourage effort input she can enlarge the delegation set by releasing the cap on top and punish the agent more if he does not acquire information by introducing a larger hole. In contrast, if it is optimal that the agent does not learn the state with probability one, then the principal would like to force the agent to take the principal-preferred no-information decision either by reducing the upper bound or enlarging the hole towards this decision.

Interestingly, we find that the principal's payoff is highest when the agent's bias is nonzero. Intuitively, the agent's bias can help the principal in the sense that the principal can punish the agent without hurting herself if the agent does not exert effort. This makes a perfectly aligned agent less appealing to the principal.

In this paper, for the sake of simplicity we assume the agent observes a perfect signal about the payoff-relevant state. One natural extension is to consider a noisy signal whose precision is increasing in the agent's effort. This would not affect our results qualitatively but more technical issues need to be taken care of. Also, we can introduce uncertainty for the principal over the agent's bias. This makes the principal's choice more subtle which is interesting to investigate. We leave this and other interesting topics to future work.

2.7 Appendix

2.7.1 Proof of Lemma 2.1

We essentially follow the proof of Theorem 1 in Holmström (1984). We endow the space of delegation sets with the Hausdorff metric, and then check that the principal is maximizing an upper semicontinuous function over a compact space.

First, select a sufficiently large compact set $Y \subseteq \mathbf{R}$ such that we may restrict to delegation sets included in Y without changing the supremum.[7] Denote by $\mathcal{K}_Y$ the space of nonempty compact subsets of Y, endowed with the Hausdorff metric. This space $\mathcal{K}_Y$ is compact (Aliprantis and Border, 2006, Theorem 3.85).

It remains to check the upper semicontinuity of the principal's objective function $U_P\colon \mathcal{K}_Y \to \mathbf{R}$ defined by

$$U_P(D) = e^*(D)u_{P,1}(D) + (1 - e^*(D))u_{P,0}(D).$$

We prove that $u_{P,1}$ and $u_{P,0}$ are upper semicontinuous and e^* is continuous.

Define the correspondence Y_0^* from $\mathcal{K}_Y$ into Y by

$$Y_0^*(D) = \operatorname*{argmax}_{y \in D} \mathbf{E}[u_A(y, \theta)].$$

Similarly, define the correspondence Y_1^* from $\mathcal{K}_Y \times [0, 1]$ into Y by

$$Y_1^*(D, \theta) = \operatorname*{argmax}_{y \in D} u_A(y, \theta).$$

Equip $\mathcal{K}_Y \times [0, 1]$ with the product topology. By Berge's theorem (Aliprantis and Border, 2006, Theorem 17.31), these correspondences are upper hemicontinuous.[8] By

[7]For example, take $Y = [-2, 3 + 2\beta]$. If a minimal delegation set is not included in this interval, then after some state realization, the agent chooses a decision at least $2 + \beta$ units away from his bliss point. So after every state realization, the agent's decision is at least $1 + \beta$ units away from his bliss point, hence at least 1 unit away from the principal's bliss point. Thus, the principal's expected loss is at least 1, which is strictly larger than the expected loss of 1/12 that the principal can secure from the singleton delegation set $\{1/2\}$.

[8]The function $y \mapsto \mathbf{E}[u_A(y, \theta)]$ is continuous by dominated convergence. For Y_0^*, we apply Berge to the correspondence φ_0 from $\mathcal{K}_Y$ into Y defined by $\varphi_0(D) = D$. The associated identity function on $\mathcal{K}_Y$ is clearly continuous, so φ_0 is continuous by Theorem 17.15 in Aliprantis and Border (2006). For Y_1^*, we apply Berge to the correspondence φ_1 from $\mathcal{K}_Y \times [0, 1]$ into Y defined by $\varphi_1(D, \theta) = D$. This correspondence is continuous because it is the composition of φ_0 with the projection map $(D, \theta) \mapsto D$.

our tie-breaking assumption, the utility functions $u_{P,0}$ and $u_{P,1}$ can be expressed as

$$u_{P,0}(D) = \max_{y \in Y_0^*(D)} \mathbf{E}[u_P(y,\theta)] \quad \text{and} \quad u_{P,1}(D) = \mathbf{E}\left[\max_{y \in Y_1^*(D,\theta)} u_P(y,\theta)\right].$$

By a variant of Berge's theorem (Aliprantis and Border, 2006, Lemma 17.30), it follows that $u_{P,1}$ and $u_{P,0}$ are upper semicontinuous.[9]

Finally, we check that e^* is continuous. Let $\Delta_A(D) = u_{A,1}(D) - u_{A,0}(D)$. It follows, as above, from Berge's theorem (Aliprantis and Border, 2006, Theorem 17.31) that Δ_A is a continuous function on $\mathcal{K}_Y$. From the agent's effort choice first-order condition, $e^*(D) = f(\Delta_A(D))$, with the function $f \colon \mathbf{R}_+ \to [0,1]$ given by

$$f(x) = \begin{cases} c'^{-1}(x) & \text{if } x < \lim_{e \to 1} c'(e), \\ 1 & \text{if } x \geq \lim_{e \to 1} c'(e). \end{cases}$$

Since c' is continuous and strictly increasing on $[0,1)$ with $c'(0) = 0$, this function f is well-defined and continuous. We conclude that $e^* = f \circ \Delta_A$ is continuous.

2.7.2 Proof of Lemma 2.2

Let D^* be a solution of (2.2). The delegation set $\{y^*(D^*)\}$ is feasible and secures the principal a payoff of $u_{P,0}(D^*)$. Since D^* is optimal, it follows that

$$u_{P,0}(D^*) \leq U_P(D^*) = e^*(D^*)u_{P,1}(D^*) + (1 - e^*(D^*))u_{P,0}(D^*),$$

so $u_{P,0}(D^*) \leq u_{P,1}(D^*)$, provided that $e^*(D^*) > 0$. If $e^*(D^*) = 0$, then minimality immediately implies that $D^* = \{y^*(D^*)\}$, so $u_{P,0}(D^*) = u_{P,1}(D^*)$. Either way, we conclude that $U_P(D^*) \leq u_{P,1}(D^*)$.

2.7.3 Proof of Theorem 2.1

Suppose $\beta \in (0, 1/2)$ and $c'(1) > \Delta_A([\beta, 1-\beta])$. Let D^* be an optimal delegation set that is minimal. If the principal offers the delegation set $[\beta, 1/2]$, the agent will choose positive effort (since $c'(0) = 0$), and therefore the principal's payoff will be

[9] For $u_{P,0}$, conclude by dominated convergence that the function $y \mapsto \mathbf{E}[u_P(y,\theta)]$ is continuous, hence upper semicontinuous. For $u_{P,1}$, Berge tells us that the integrand is upper semicontinuous in (D,θ), so $u_{P,1}$ is upper semicontinuous by dominated convergence.

strictly greater than the payoff $-1/12$ from the singleton delegation set $\{1/2\}$. We conclude that

$$u_{P,1}(D^*) > -1/12 \geq u_{P,0}(D^*) \quad \text{and} \quad e^*(D^*) > 0.$$

Since D^* is compact, it has a minimum and a maximum. First we check that the minimum is at most β.

Lemma 2.3 (Minimal decision)**.** We have $\min D^* \leq \beta$.

Proof of Lemma 2.3. Let $\underline{d} = \min D^*$. We must have $\underline{d} < 1 + \beta$, for otherwise every decision is at least $1 + \beta$, which cannot be optimal. Suppose for a contradiction that $\underline{d} > \beta$. Append the point β to the delegation set D^*. Since $\underline{d} < 1+\beta$, the uninformed action does not change. The informed action does change, and $u_{A,1}$ strictly increases. Moreover, the mean of the informed action strictly decreases, so $u_{P,1}$ must strictly increase. Therefore, Δ_P and Δ_A strictly increase. We conclude that U_P strictly increases, contrary to the optimality of D^*. □

The next lemma says that below $1 + \beta$, the only hole that D^* can have is a symmetric hole around $1/2 + \beta$.

Lemma 2.4 (Holes)**.** Fix decisions d_1 and d_2 satisfying $d_1 < d_2$ and $d_2 \in [\beta, 1 + \beta]$. If $[d_1, d_2] \cap D^* = \{d_1, d_2\}$, then $(d_1 + d_2)/2 = 1/2 + \beta$.

Proof of Lemma 2.4. We prove this result by contradiction. There are two cases.

First suppose $1/2 + \beta \notin (d_1, d_2)$. In particular, this implies $(d_1 + d_2)/2 \neq 1/2 + \beta$. "Filling in the hole," i.e., appending (d_1, d_2) to D^* leaves the uninformed action unchanged and strictly decreases the variance of the informed action, without increasing its mean.[10] Therefore, $u_{P,1}$ and $u_{A,1}$ both strictly increase (while $u_{P,0}$ and $u_{A,0}$ are unchanged). Hence U_P strictly increases, contrary to optimality.

Next suppose $1/2 + \beta \in (d_1, d_2)$ and $(d_1 + d_2)/2 \neq 1/2 + \beta$. Thus, one endpoint of (d_1, d_2) is strictly closer to $1/2 + \beta$ than the other. We split into two cases.

If d_1 is closer, append $[1+2\beta-d_1, d_2)$ to D^*. If d_2 is closer, append $(d_1, 1+2\beta-d_2]$ to D^*. Either way, the variance of the informed action strictly decreases while the mean

[10]This is where we use the fact that $d_2 \leq 1 + \beta$. It is clear that the variance decreases. If $d_1 \geq \beta$, then the mean is unchanged. Otherwise, the mean strictly decreases. This can be seen geometrically by extending the state space to $d_1-\beta$ and observing that the mean would be unchanged on this extended state space. Since the mean increases on $[d_1 - \beta, 0]$ and remains unchanged over $[d_1 - \beta, d_2 - \beta]$, it must decrease on $[0, d_2 - \beta]$ as claimed.

is unchanged. The uninformed action can change, but the agent's uninformed payoff $u_{A,0}$ does not. Because of our tie-breaking convention for the uninformed action, the principal's uninformed payoff $u_{P,0}$ can only increase, as we now check. If d_2 is closer to $1/2+\beta$, then the uninformed action changes from d_2 to $1+2\beta-d_2$, which is closer to $1/2$. If d_1 is closer, the uninformed action is unchanged. We conclude that in both cases, $u_{P,1}$ and $u_{A,1}$ strictly increase, $u_{A,0}$ is unchanged, and $u_{P,0}$ does not decrease. Therefore, the principal's payoff U_P strictly increases, contrary to optimality. □

With these lemmas established, we now complete the proof, separating into cases according to the value of $\max D^*$.

Case 1 (Interval delegation). Suppose $\max D^* \leq 1/2 + \beta$. Set $\bar{y} = \max D^*$. By Lemma 2.4, holes must be centered on $1/2 + \beta$, which is impossible in this case. We conclude that there are no holes, so minimality implies $D^* = [\beta, \bar{y}]$.

It remains to prove that $\bar{y} > 1/2$. Clearly $\bar{y} \neq \beta$, for otherwise the principal could do strictly better by offering the singleton delegation set $\{1/2\}$. So suppose for a contradiction that $\bar{y} \in (\beta, 1/2]$. Let $U_P(y) = U_P([\beta, y])$ for $y \geq \beta$, and similarly parametrize all functions of the delegation set via the map $y \mapsto [\beta, y]$. Optimality implies that $\Delta_P(\bar{y}) \geq 0$, and it is straightforward to check that $e^{*\prime}(\bar{y}) \geq 0$ and $u'_{P,0}(\bar{y}) \geq 0$. Therefore,

$$U_P{}'(\bar{y}) = e^{*\prime}(\bar{y})\Delta_P(\bar{y}) + e^*(\bar{y})u'_{P,1}(\bar{y}) + (1 - e^*(\bar{y}))u'_{P,0}(\bar{y}) \geq e^*(\bar{y})u'_{P,1}(\bar{y}).$$

The right hand side of the inequality is strictly positive because $e^*(\bar{y}) > 0$ (since $\bar{y} > \beta$) and $u'_{P,1}(\bar{y}) > 0$ (since $\bar{y} < 1 - \beta$).

Case 2 (Hollow delegation). Suppose $1/2 + \beta < \max D^* \leq 1 + \beta$. Set $\bar{y} = \max D^*$. By Lemma 2.4, holes must be centered on $1/2 + \beta$, so minimality implies $D^* = [\beta, \bar{y}] \setminus B_r(1/2 + \beta)$ for some radius $r \in [0, \bar{y} - 1/2 - \beta]$.

It remains to prove that $\bar{y} \in (1 - \beta, 1 + \beta)$ and $r > 0$. We prove these in turn. Taking r as fixed, let $U_P(y) = U_P([\beta, y] \setminus B_r(1/2+\beta))$ for $y \geq 1/2+\beta+r$, and similarly parametrize all functions of the delegation set via the map $y \mapsto [\beta, y] \setminus B_r(1/2 + \beta)$. Since the cap does not change the uninformed decision, we know $u'_{P,0}(\bar{y}) = 0$, so

$$U'_P(\bar{y}) = e^{*\prime}(\bar{y})\Delta_P(\bar{y}) + e^*(\bar{y})u'_{P,1}(\bar{y}). \tag{2.7.4}$$

Since the principal can secure a payoff strictly greater than $-1/12$ by offering the delegation set $[\beta, 1/2]$, optimality implies that $\Delta_P(\bar{y}) > 0$ and $e^*(\bar{y}) > 0$.

Suppose for a contradiction that $\bar{y} \leq 1-\beta$. Then $e^{*\prime}(\bar{y}) \geq 0$ with strict inequality if $e^*(\bar{y}) < 1$, and $u'_{P,1}(\bar{y}) \geq 0$ with strict inequality if $\bar{y} < 1-\beta$. This gives the desired contradiction $U'_P(\bar{y}) > 0$ unless $e^*(\bar{y}) = 1$ and $\bar{y} = 1-\beta$. In this last case, we must have $r > 0$ because $c'(1) > \Delta_A([\beta, 1-\beta])$. But then, since $u'_{P,1}(1-\beta) = 0$ and $u'_{A,1}(1-\beta) > 0$, the principal can strictly increase her payoff by slightly decreasing the radius and slightly increasing the cap, while maintaining full information acquisition.

Next, suppose for a contradiction that $\bar{y} = 1+\beta$. The delegation set $\{\beta, 1+\beta\}$ can be ruled out by direct computation,[11] so we may assume $r < 1/2$. To get a contradiction, we prove that U'_P is strictly increasing at $\bar{y} = 1+\beta$. Note that $u'_{P,1}(1+\beta) = 0$, and also $e^{*\prime}(1+\beta) = 0$ because $u'_{A,1}(1+\beta) = 0$. Hence, (2.7.4) gives $U'_P(1+\beta) = 0$, so we turn to the second derivatives. Since the first derivatives vanish, we get

$$U''_P(1+\beta) = e^{*\prime\prime}(1+\beta)\Delta_P(1+\beta) + e^*(1+\beta)u''_{P,1}(1+\beta).$$

It can be checked that $u''_{A,1}(1+\beta) = 0$, so $e^{*\prime\prime}(1+\beta) = 0$ and the first term vanishes. The second term is strictly positive because $u''_{P,1}(1+\beta) = 2\beta > 0$. Therefore, $U''_P(1+\beta)$ is strictly positive, which gives the contradiction.

Having proved that $\bar{y} \in (1-\beta, 1+\beta)$, we now check that r is nonzero. Suppose for a contradiction that $r = 0$. We claim that the principal can strictly increase her payoff by slightly increasing the radius and slightly decreasing the cap, while keeping Δ_A constant. Observe that increasing the radius causes a first-order increase in $u_{P,0}$, a second-order decrease in $u_{A,0}$, and third-order decreases in $u_{P,1}$ and $u_{A,1}$. Decreasing the cap will cause a first-order increase in $u_{P,1}$ and a first-order decrease in $u_{A,1}$. It follows that for sufficiently small changes, both $u_{P,0}$ and $u_{P,1}$ will strictly increase, yielding the contradiction.

[11] If $\beta \in [1/4, 1/2)$, the delegation set $\{\beta, 1+\beta\}$ results in a loss strictly larger than $1/12$. If $\beta \in (0, 1/4)$, the delegation set $\{\beta, 1+\beta\}$ results in at least a loss of

$$\begin{aligned}(1/3)((1/2+\beta)^3 - \beta^3) + (1/3)(\beta^3 + (1/2-\beta)^3) &= (1/3)((1/2+\beta)^3 + (1/2-\beta)^3) \\ &> (1/3)2(1/2)^3 = 1/12,\end{aligned}$$

where the strict inequality follows from Jensen's inequality.

Case 3 (High-point delegation). Suppose $\max D^* > 1 + \beta$. Set $h = \max D^*$. Let $y_0 = \max(D^* \setminus \{h\})$. If $y_0 \leq \beta$, then we can argue directly that the principal's utility is strictly less than $-1/12$.[12] Hence $y_0 > \beta$.

We will prove that $h < 2 + 2\beta - y_0$ and $h > 1 + y_0$. In the process, we will prove that $y_0 < 1/2 + \beta$, so it then follows from the lemmas that there can be no holes, hence $D^* = [\beta, y_0] \cup \{h\}$.

Since D^* is minimal, the agent must strictly prefer h to y_0 when $\theta = 1$, so $(h + y_0)/2 < 1 + \beta$, hence $h < 2 + 2\beta - y_0$. Next, observe that appending $2 + 2\beta - h$ (and removing h) increases both players' informed payoffs, so optimality requires that it change the uninformed decision. In particular, this means $y_0 < 1/2 + \beta$. Moreover, $2 + 2\beta - h < 1 + 2\beta - y_0$, so $h > 1 + y_0$.

2.7.4 Proof of Theorem 2.2

For the comparative statics in the cost, fix the bias β, and suppose the cost function satisfies $c'(1) > \Delta_A([\beta, 1 - \beta])$. Let D^* be an optimal delegation set.

If the agent does not choose full effort, then reducing the marginal cost pointwise will strictly increase effort, and hence strictly increase the principal's payoff from the set D^*.

If the agent chooses full effort, then we separate into cases. If D^* is an interval, then full acquisition implies that $\bar{y} > 1 - \beta$. If the marginal cost decreases pointwise, then the principal can strictly improve her payoff by slightly reducing the cap $\bar{y}$, while maintaining full acquisition. If D^* is hollow, then the principal can strictly improve her payoff by slightly reducing the radius r, while maintaining full acquisition. Finally, if D^* has a high point, then the principal can slightly increase y_0, while maintaining full acquisition. This strictly increases the principal's payoff because $\partial u_{P,1}/\partial y_0 = (1/4)(h - y_0)^2 > 0$.

For the comparative statics in the bias, fix a cost function c such that it is optimal to induce an unbiased agent to acquire only partial information. The optimal delegation set is $[0, 1] \setminus B_r(1/2)$ for some positive radius r. Let e_0 be the (unbiased) agent's effort choice when offered the delegation set $[0, 1] \setminus B_r(1/2)$. By assumption, $e_0 < 1$.

[12] Suppose $D = \{y_0\} \cup \{h\}$ for $y_0 \leq \beta < 1 + \beta < h$. It is easy to show that $u_{P,1}(D) \leq u_{P,1}(\{\beta\} \cup \{h\})$. For $h \in (1 + \beta, 2 + \beta)$, the utility $u_{P,1}(\{\beta\} \cup \{h\})$ has a single local maximum, so it suffices to check the endpoints $h = 1 + \beta$ and $h = 2 + \beta$. In both cases, the utility $u_{P,1}(\{\beta\} \cup \{h\})$ is strictly less than $-1/12$.

For each β, let $D(\beta) = [\beta, 1+\beta] \setminus B_r(1/2+\beta)$. If the principal offers the delegation set $D(\beta)$ to an agent with bias β, then the agent will still choose effort level e_0.

For each bias β, let $V(\beta)$ denote the principal's expected payoff from optimal delegation. Let $U_P(D;\beta)$ denote the principal's expected payoff from offering delegation set D when the agent's bias is β. We have

$$\begin{aligned} V(\beta) - V(0) &\geq U_P(D(\beta);\beta) - U_P(D(0);0) \\ &= -e_0\beta^2 + (1-e_0)(r^2 - (r-\beta)^2). \end{aligned}$$

The right side vanishes at $\beta = 0$, and its derivative is

$$-2\beta e_0 + 2(1-e_0)(r-\beta),$$

which is strictly positive for $\beta < (1-e_0)r$. So in fact we have proven that the the principal's payoff is increasing in the agent's bias over the interval $[0, (1-e_0)r]$, which is nontrivial because $e_0 < 1$.

2.7.5 Proof of Proposition 2.1

Suppose $\kappa > \Delta_A([\beta, 1-\beta])$.

Let D^* be an optimal delegation set, and suppose $e^*(D^*) = 1$. We prove the three parts of the proposition separately.

Interval Suppose $D^* = [\beta, \bar{y}]$ for some $\bar{y} \in (1/2, 1/2+\beta]$. The agent's condition for the full acquisition gives

$$\kappa \leq \Delta_A(D^*) = (1/3)(\bar{y}-\beta)^3,$$

hence

$$\bar{y} \geq \beta + (3\kappa)^{1/3}.$$

Since $\kappa > \Delta_A([\beta, 1-\beta])$, full acquisition implies that $\bar{y} > 1-\beta$. Therefore, optimality implies that the inequality $\bar{y} \geq \beta + (3\kappa)^{1/3}$ holds with equality (for otherwise the principal could do strictly better by slightly reducing the cap).

Hollow with upper point Suppose $D^* = [\beta, \bar{y}] \setminus B_r(1/2+\beta)$ and moreover that $\bar{y} = 1/2 + \beta + r$. Hence $r \in (0, 1/2)$. Thus,

$$D^* = [\beta, 1/2 + \beta - r] \cup \{1/2 + \beta + r\}.$$

We will show that in this case, the radius r is independent of β and strictly increasing in κ. Note that

$$\Delta_A(D^*) = (1/24)(1 + 2r(3 + 6r - 4r^2)).$$

The agent's full acquisition condition gives $\kappa \leq \Delta_A(D^*)$, and this must hold with equality, for otherwise the principal could strictly increase her payoff by strictly decreasing the radius. We have

$$\kappa = (1/24)(1 + 2r(3 + 6r - 4r^2)),$$

which defines an implicit function $r(\kappa)$. Differentiating with respect to κ gives

$$\frac{\partial r}{\partial \kappa} = \frac{1}{1/4 + r - r^2} > 0.$$

Therefore, r is strictly increasing in κ and independent of β. The claimed comparative statics follow.

Hollow with upper interval Suppose

$$D^* = [\beta, \bar{y}] \setminus B_r(1/2 + \beta),$$

for some $\bar{y} \in (\max\{1/2+\beta, 1-\beta\}, 1+\beta)$ and some $r \in (0, \bar{y} - 1/2 - \beta]$. In particular, we know that the solution must maximize $u_{P,1}$ subject to the constraint $\kappa = \Delta_A$. The first-order conditions of this problem give

$$\bar{y} = 1 - \beta + 2r\beta.$$

Plugging this back into Δ_A, we obtain

$$\Delta_A = 1/12 + r^2 - (2r^3)/3 + 8/3(-1 + r)^3\beta^3.$$

Hence,

$$\kappa = 1/12 + r^2 - (2r^3)/3 + 8/3(-1+r)^3\beta^3,$$

which defines an implicit function $r(\kappa, \beta)$. Take derivatives to get

$$\frac{\partial r}{\partial \kappa} = \frac{1}{2r(1-r) + 8(1-r)^2\beta^3} > 0,$$
$$\frac{\partial r}{\partial \beta} = \frac{8(1-r)^3\beta^2}{2r(1-r) + 8(1-r)^2\beta^3} > 0.$$

It follows that r is strictly increasing in κ and β.

2.7.6 Proof of Proposition 2.2

Suppose $\kappa > \Delta_A([\beta, 1-\beta])$. Let D^* be an optimal delegation set, and suppose $e^*(D^*) < 1$. We prove the two parts of the proposition separately.

Interval Suppose $D^* = [\beta, \bar{y}]$ for some $\bar{y} \in (1/2, 1/2+\beta]$. Write $U_P(\bar{y})$ for $U_P([\beta, \bar{y}])$. We have

$$U_P = -(1/3) + \bar{y} - \bar{y}^2 + (9\kappa)^{-1}(\bar{y}-\beta)^5(\bar{y}+2\beta),$$
$$\frac{\partial U_P}{\partial \bar{y}} = 1 - 2\bar{y} + (3\kappa)^{-1}(\bar{y}-\beta)^4(2\bar{y}+3\beta),$$
$$\frac{\partial^2 U_P}{\partial \bar{y}^2} = -2 + (10/3)\kappa^{-1}(\bar{y}-\beta)^3(\bar{y}+\beta),$$
$$\frac{\partial^3 U_P}{\partial \bar{y}^3} = (20/3)\kappa^{-1}(\bar{y}-\beta)^2(2\bar{y}+\beta).$$

Uniqueness of optimal $\bar{y}$ follows from the third derivative which is strictly positive.

For the comparative statics, observe that

$$\frac{\partial^2 U_P}{\partial \bar{y}\partial \beta} = -(5/3)\kappa^{-1}(\bar{y}-\beta)^3(\bar{y}+3\beta),$$
$$\frac{\partial^2 U_P}{\partial \bar{y}\partial \kappa} = -(1/3)\kappa^{-2}(\bar{y}-\beta)^4(2\bar{y}+3\beta).$$

Both these expressions are strictly negative.

It is clear that the optimal $\bar{y}$ converges to $1/2$ as $\beta \to 1/2$ or $\kappa \to \infty$ from the first-order condition, $\frac{\partial U_P}{\partial y} = 0$.

Finally, we check that $\bar{y} < 1-\beta$. To see this, we will show that for $\bar{y} \geq 1-\beta$,

$$\frac{\partial U_P}{\partial \bar{y}} \geq 0 \implies \frac{\partial^2 U_P}{\partial \bar{y}^2} > 0.$$

Indeed, it suffices to show that

$$(3\kappa)^{-1}(\bar{y}-\beta)^3 > 1/5,$$

and this follows from the fact that

$$(3\kappa)^{-1}(\bar{y}-\beta)^3 \geq \frac{2\bar{y}-1}{(\bar{y}-\beta)(2\bar{y}+3\beta)} \geq \frac{1}{2\bar{y}+3\beta} \geq \frac{1}{2+5\beta} > 1/5.$$

Hollow The principal's utility from a hollow delegation set $[\beta, \bar{y}] \setminus B_r(1/2+\beta)$ is given by

$$U_P = u_{P,0} + \frac{\Delta_A \Delta_P}{\kappa}.$$

Now we view each expression as a function of r and let $f(r) = \Delta_A(r)\Delta_P(r)$. Fix $\kappa > \kappa'$, and let r be an optimal radius at κ and let r' be an optimal radius at κ'. Optimality implies that

$$\begin{aligned} u_{P,0}(r) + f(r)/\kappa &\geq u_{P,0}(r') + f(r')/\kappa, \\ u_{P,0}(r') + f(r')/\kappa' &\geq u_{P,0}(r) + f(r)/\kappa'. \end{aligned}$$

Adding these inequalities and simplifying show that $f(r') \geq f(r)$, hence $u_{P,0}(r) \geq u_{P,0}(r')$, so $|r-\beta| \leq |r'-\beta|$. From the first-order condition, we can see that $r \neq r'$, so this inequality must be strict.

The final observation follows from the fact that Δ_A is bounded over optimal delegation sets, and hence $\Delta_A/\kappa \to 0$ as $\kappa \to \infty$.

2.7.7 Example: Optimality of High-point Delegation

The idea is to select a cost function so that the effort best response e^* is very sensitive near a particular threshold and that the effort choice near this threshold is small. Then the optimal delegation set must induce an uninformed action that is very near 1/2 but also provide enough flexibility so that the agent's return to information crosses this threshold. A high-point delegation set will have the desired properties.

We begin with a technical lemma.

Lemma 2.5. For $\beta \in (1/4, 1/2)$, there exists a high point $h \in (3/2, 3/2 + 2\beta)$ and a radius $\rho \in (0, \beta)$ satisfying

$$\begin{aligned}
\Delta_A([\beta, 1/2] \cup \{h\}) &> \Delta_A([\beta, 1/2 + \rho]), \\
\Delta_P([\beta, 1/2] \cup \{h\}) &> 0, \\
u_{P,1}([\beta, 1/2] \cup \{h\}) &> u_{P,1}([\beta, 1/2 + \rho] \cup \{1/2 + 2\beta - \rho\}).
\end{aligned}$$

Proof of Lemma 2.5. Fix $\beta \in (1/4, 1/2)$. Then this follows from observing that all the payoffs from $[\beta, 1/2] \cup \{h\}$ converge to the corresponding payoffs from $[\beta, 1/2]$ as $h \uparrow 3/2 + 2\beta$. We have

$$\Delta_P([\beta, 1/2]) = (1/24)(1 - 2\beta)^2(1 + 4\beta) > 0$$

and

$$u_{P,1}([\beta, 1/2]) - u_{P,1}([\beta, 1/2] \cup \{1/2 + 2\beta\}) = (\beta/2)(4\beta - 1) > 0.$$

Therefore, the last inequality in the lemma holds as long as h and ρ are sufficiently small. Fix such a ρ and then pick h sufficiently small to satisfy the top two inequalities in the lemma. □

Fix $\beta \in (1/4, 1/2)$. To complete the example, we will find a cost function under which $[\beta, 1/2] \cup \{h\}$ yields higher utility to the principal than all interval and hollow delegation sets.

First we construct the cost function. Denote the left and right hand side of the first inequality in Lemma 2.5 by δ_1 and δ_0. For each $\varepsilon \in (0, 1/2)$, we may select a cost function c such that[13]

$$\delta_0 < c'(\varepsilon^2) < c'(\varepsilon - \varepsilon^2) < \delta_1 < 1 < c'(\varepsilon).$$

Set $D_h = [\beta, 1/2] \cup \{h\}$ and $D_\rho = [\beta, 1/2 + \rho] \cup \{1/2 + 2\beta - \rho\}$. First, note that

$$U_P(D_h) \geq -1/12 + (\varepsilon - \varepsilon^2)\Delta_P(D_h),$$

[13] It is straightforward to construct a continuously differentiable derivative function with strictly positive derivative that satisfies these bounds and also the equality $c'(0) = 0$. To get the cost function itself, integrate this derivative function from 0.

where $\Delta_P(D_h)$ is positive by Lemma 2.5. We want to show that the expression on the right hand side of this inequality is an upper bound on the principal's utility from interval and hollow delegation sets. Let D be such a delegation set. There are three cases.

1. If $u_{P,0}(D) < -1/12 - \rho^2$, then

$$U_P(D) \leq (1 - e^*(D))u_{P,0}(D) < -(1-\varepsilon)(1/12 + \rho^2).$$

2. If $D = [\beta, \bar{y}]$ for some $\bar{y} \leq 1/2 + \rho$, then $\Delta_A(D) \leq \Delta_A([\beta, 1/2 + \rho]) = \delta_0$, so by the construction of the cost function, $e^*(D) \leq \varepsilon^2$. Since $\Delta_P(D) \leq 1$, we have

$$U_P(D) \leq -1/12 + e^*(D) \leq -1/12 + \varepsilon^2.$$

3. Suppose $D = [\beta, \bar{y}] \setminus B_r(1/2 + \beta)$ for some cap $\bar{y}$ and radius $r \in [\beta - \rho, \beta + \rho]$. Compare D to D_ρ, and observe that D_ρ has a smaller radius and cap closer to $1 - \beta$ (since $\rho \leq \beta$, as guaranteed by Lemma 2.5). Hence $u_{P,1}(D) \leq u_{P,1}(D_\rho)$. Therefore,

$$\begin{aligned} U_P(D) &\leq -1/12 + e^*(D)(u_{P,1}(D) + 1/12) \\ &\leq -1/12 + \varepsilon(u_{P,1}(D) + 1/12)_+ \\ &\leq -1/12 + \varepsilon(u_{P,1}(D_\rho) + 1/12)_+, \end{aligned}$$

   where we have used the fact that $\Delta_A(D) \leq 1$, so $e^*(D) \leq \varepsilon$. By Lemma 2.5, $(u_{P,1}(D_\rho) + 1/12)_+ < \Delta_P(D_h)$.

In each case, provided that ε is sufficiently small, our upper bound on $U_P(D)$ is strictly below our lower bound on $U_P(D_h)$. Choose ε small enough such that all three cases go through.

Chapter 3

Reputational Cheap Talk vs. Reputational Delegation

3.1 Introduction

Information is crucial for decision makings. Yet it is dispersed among different parties. It is quite often that the decision maker (henceforth, principal) does not have the relevant information and therefore needs to rely on the help by an expert (henceforth, agent). Investors, for example, are often less informed about the profitability of a project than the managers. Likewise, a monopolist is often better informed about his costs than the regulator. To benefit from the agent's information, the principal can either ask him to report what he knows (communication), or she can simply let the agent make the decision himself (delegation). However, the agent behaves in accordance with his own preferences, which are not necessarily aligned with the principal's. Therefore, the principal is concerned with the information distortion if communicating with the agent, or extreme actions the agent may take if delegating the authority to the agent.

If the agent's preferences are common knowledge, Dessein (2002) resolves the trade-off between the two decision making protocols and shows that delegation dominates communication as long as the agent's bias is not too large. However, when the principal is uncertain about the agent's preferences, at first sight it is not so straightforward that whether this result still applies. Compared to the no uncertainty case, now different types of the agent may behave differently: The good agent may want to separate from the bad agent, which might reverse the dominance relation between the two decision making protocols.

Beside uncertainty over the agent's preferences, it is also natural to consider repeated interactions between the principal and the agent, which are pervasive in practice. In a long run relationship, the principal obtains one more device to discipline the agent: The agent wants to "look good," i.e., has reputational concerns. It is well known in economic theory that the reputation of an agent provides him with implicit commitment power and can thereby substitute for explicit contractual enforcement.

In this paper, I address the question that whether the principal should keep the authority of decision making when the agent has reputational concerns.

I study a simple two-period model with two states and continuous actions. Only the agent knows the payoff-relevant state in each period. The principal is not sure about the agent's type. The agent can be perfectly congruent with the principal (good agent) or he can be biased (bad agent) so that he always prefers a higher action no matter the state compared to the principal. In each period of the game, the principal first decides whether to delegate the decision making to the agent. If she does, the agent takes an action himself. If she does not, the agent sends a cheap talk message to the principal who then takes an action.

In the second period there is no reputational concerns. It can be seen as the static benchmark, i.e., the one-period version of the dynamic model. Since this is the last stage of the game, the agent will choose his most preferred action if he has the authority to make a decision. Whereas, if the principal keeps control, there exists an unique informative equilibrium in which the good agent truthfully reports the state while the bad agent always claims that he observes a high state. Intuitively, the good agent always strongly prefers to reveal the true state since he has the same preferences with the principal. Given the binary structure of the state (i.e., the states are sufficiently far apart), the bad agent has no incentive to mimic the good agent's report when seeing a low state given his strong preferences over the high action of the principal. In this equilibrium, the principal's expected utility is only determined by the prior belief about the agent's type at the beginning of the second period.

Given the structure of this unique informative communication equilibrium, I find that the principal is better off by keeping control instead of delegating to the agent. This result contrasts with the one obtained by Dessein (2002) who shows that delegation dominates communication whenever informative communication equilibrium is available. As he points out, the key to his analysis is that the agent's bias is systematic and predictable. If the agent's preferences are uncertain to the principal, noisy communication can be optimal for the principal when she is too afraid of the agent taking extreme actions.

Interestingly, the first-period game can be transformed into a costly signaling game in the spirit of Spence (1973). The agent's behavior not only (in)directly affects his current utility but also signals his type which affects his future utility. If the principal keeps control, the unique informative equilibrium features the same structures as that of the second period communication game. Although the agent has reputational considerations, he cares about the two periods equally, and thus the reputation gain by mimicking the good agent is not sufficient to compensate the current utility loss. Consequently, while the good agent chooses to tell the truth, the bad agent always announces he observes the high state.

If the principal delegates control, given the continuous action space, multiple equilibria can arise. In particular, the bad agent's action can also be sensitive to the state. With full discretion to choose an action, a lower current utility can be compensated by a higher reputation gain so that the bad agent's action is not independent of the state any more. This provides possibilities of delegation improving upon communication when the agent has reputational concerns.

Specifically, I find that in the first period, if the fraction of good agent is relatively high, delegation dominates communication while this relation is overturned when the fraction of good agent is relatively low. Intuitively, when the principal believes the agent is more likely to be aligned, it is less costly for the good agent to signal his type for a higher reputation in future. Also delegation makes the agent with authority behave more sensitively to changes in the state when making decisions. And thus when the good agent is sufficiently populated, the principal would like to let the agent make decisions to avoid the loss of information. However, when the principal is pessimistic about the agent's motives, the good agent needs to behave more aggressively to separate from the bad agent. This reputation effect makes the principal suffer more from the loss of control so that she would rather prefer to ask the agent for information and keep control.

Related Literature

This paper is motivated by the comparison between communication and delegation initiated by Dessein (2002) in a static setting.[1] I generalize this comparison to a dynamic setting with the assumption that the principal is uncertain about the agent's preferences so that the agent's reputational concerns arise naturally. Therefore, this paper belongs to the literature on reputation and career concerns, where the agent's past actions determine his future opportunities and payoffs.

In my model, the agent is either engaged in a cheap talk communication game or a delegation problem. In the communication game, the agent's reputation is built upon the messages he sends. These messages signal his preferences. The setting mostly follows Morris (2001), who extends the repeated cheap talk model with reputation studied in Sobel (1985) and Benabou and Laroque (1992) to a setting where the good agent does not commit to tell the truth. Compared to Morris (2001), I assume the principal cannot observe the state throughout the game while the state is publicly observed at the end of each period in Morris (2001). Another difference in modeling is that I assume the agent perfectly observes the state rather than a noisy signal of the state as in Morris (2001), which simplifies the analysis without losing qualitative generality.

The framework that I follow to build reputational concerns into the model is different from the reputational cheap talk model studied by Ottaviani and Sørensen (2006a,b) (and similarly, Scharfstein and Stein (1990), Prendergast and Stole (1996), Levy (2004), Prat (2005), and Gentzkow and Shapiro (2006)), where the expert is signaling his expertise rather than preferences, i.e., the agent is concerned to look well informed rather than aligned with the principal.

Another related paper is Avery and Meyer (2012) who consider a two-period and two-action model with a potentially biased agent. The bias may be either low or high. In the baseline model where the agent's bias and the distribution of the state are constant across periods, they show that the principal benefits from the reputational incentives whereas this result may be overturned in a more general model where either

[1] Ottaviani (2000) also compares communication with delegation in a static setting. He discusses this issue in a number of variants of the quadratic-uniform model which is the leading example in Crawford and Sobel (1982). The variations include noisy information, uncertain bias and possibly naive principal, etc. Lately, Deimen and Szalay (2019) investigate the performance of communication relative to delegation provided that the principal relies on the agent to acquire information so that the agent's information is endogenous rather than exogenously given.

the agent's bias or the state distribution varies across periods. Unlike my paper, they mainly address the question that whether the principal is better off in a repeated cheap talk game compared to a one-period static benchmark where the agent has no reputational concerns.

Other work in this area has focused on different situations. Morgan and Stocken (2003) also consider a cheap talk game where there is uncertainty about the incentives of the agent. Since the setting is static, they essentially extend the seminal work of Crawford and Sobel (1982) in which the agent is surely biased to the case where the agent is potentially biased. As another extension of Crawford and Sobel (1982), Golosov et al. (2014) study a finite-horizon dynamic strategic information transmission problem with no uncertainty about the agent's preferences and therefore reputational concerns are absent. Xu (2012) constructs a reputational cheap talk model with information acquisition and mainly focuses on the effect of reputational concerns on the agent's information acquisition incentives. He also shows that delegation may reduce the good agent's information acquisition incentives and due to this effect communication is better for the principal than delegation.

Other than cheap talk, this paper is also related to the delegation literature.[2] In particular, it contributes to the growing literature on dynamic delegation. Most closely related is Ely and Välimäki (2003) who construct a model where short-run uninformed principals decide whether to hire a long-run informed agent at each period. The principals do not know the agent's type and they aim to hire the good agent and fire the bad one. In each period, the agent, if hired, takes a payoff-relevant action to signal his type like what the agent does in my model. However, they focus on the distortionary effects of the incentive of avoiding "looking bad" but I mainly investigate whether delegation dominates communication when the agent has reputational considerations.

Lipnowski and Ramos (2018) also consider an infinite-horizon repeated game in which at each period an uninformed principal chooses whether to delegate the project adoption choice to an informed agent who knows the current project is good or bad. They show that at the early stages the principal always lets the agent make the decision and as the average quality of the adopted projects drops to some level, the

[2] In this paper if the principal delegates control to the agent, the agent can choose whatever decision he prefers. This full delegation is different from the constrained delegation problem initiated by Holmström (1977, 1984).

principal would delegate less than before. Unlike my model, Lipnowski and Ramos (2018) assume the agent's bias is commonly known and hence there is no reputational considerations, although we share the feature that the principal has no commitment power and the delegation decision only depends on the agent's past actions. In contrast, Guo and Hörner (2015) assume the principal can commit to a specific decision rule ex ante and they fully characterize the optimal contract in a dynamic setting where an uninformed principal decides whether to provide a costly and perishable good to the agent at each period.[3]

The remainder of this paper is organized as follows. Section 3.2 presents the model. Section 3.3 first characterizes the equilibrium under each decision making protocol, and then analyzes the optimal authority allocation of the second period of the game. Section 3.4 analyzes the first period of the game with reputational incentives operative and proceeds as Section 3.3. Section 3.5 concludes.

3.2 Model

There are two periods of interactions between a principal (she, P) and an agent (he, A). They both know that the period-t state of the world θ_t is equally likely to be either 0 or 1 and that θ_1 and θ_2 are independent. The agent privately knows his utility function. In contrast, the principal is uncertain about the agent's preferences. At the beginning of the first period, it is a common prior that the agent has the same utility function with the principal with probability π_1 (good agent, $A =$ "G") and prefers a higher action at each state with probability $1 - \pi_1$ (bad agent, $A =$ "B"), where $0 < \pi_1 < 1$. Specifically, the preferences of the principal and the good agent are characterized by the following quadratic-loss utility functions which depend on the action taken by either the principal or the agent and the state of the world:

$$U_t^i(a_t, \theta_t) = -(a_t - \theta_t)^2 \quad \text{for } i = P, G \text{ and } t = 1, 2$$

and the payoff to the bad agent is given by

$$U_t^B(a_t, \theta_t) = a_t \quad \text{for } t = 1, 2$$

[3]For other works on dynamic delegation, see for example, Alonso and Matouschek (2007), Guo (2016), Bird and Frug (2017), and Li et al. (2017).

which is taken to be the action in each period for simplicity. That is, the principal's flow payoff is maximized by matching the action with the state, the good agent is perfectly aligned with the principal, but there is a conflict of interest between the bad agent and the principal.

The two periods are equally important for both the principal and the agent, i.e., no discounting. The principal and the agent's aggregate utility is given by $U_1^P + U_2^P$ and $U_1^A + U_2^A$, respectively.

I assume that the state is not publicly observable throughout the game and only the agent is perfectly informed about the state in each period. In particular, the principal does not observe her first-period payoff before choosing her second-period authority allocation. I restrict attention to pure strategy equilibria.

The game proceeds as follows. (i) At the beginning of the first period, the principal decides whether or not to delegate the decision making to the agent. (ii) The agent observes the state of the world θ_1. If the principal keeps the authority, then the agent sends a costless and unverifiable message to the principal. Given the agent's message, the principal takes an action. If the principal delegates control to the agent, then the agent takes an action himself. The second period of the game repeats the first period with a new and independent state θ_2.

3.3 Second Period of the Game

I start the analysis with the second period. Since this is the last stage of the game, the agent simply maximizes his current utility without any reputational concern. Suppose at the beginning of the second period, the principal believes that the probability of the agent being good is π_2, which is also referred to as the reputation of the agent. Clearly, it is influenced by the message sent by the agent if there was communication in the first period, or the action taken by the agent if the agent was awarded the authority in the first period. The agent's reputational concern arises because it will influence the second period's play.

3.3.1 Delegation

If the principal delegates the decision making to the agent, the good agent simply implements $a_2^G = \theta_2$ and the bad agent implements $a_2^B = 1$ regardless of the state.

The principal's expected payoff from delegating control is

$$\mathbb{E}U_2^P = \frac{1}{2}(1-\pi_2)\cdot[-(1-0)^2] = -\frac{1-\pi_2}{2}. \tag{1}$$

Note that the principal will find it worthwhile to delegate control only if it yields her weakly higher payoff than she keeps control and takes an uninformed action.[4] The principal's best uninformed decision is $\mathbb{E}\theta_2 = 1/2$ which yields her expected utility $-1/4$.[5] Hence, the principal will delegate authority only if

$$-\frac{1-\pi_2}{2} \geq -\frac{1}{4},$$

or

$$\pi_2 \geq \frac{1}{2}. \tag{2}$$

3.3.2 Communication

I first introduce some notations. Given that I focus on pure strategy equilibria, let $\sigma_A : \{0,1\} \to M$ be the agent's reporting strategy which assigns a message of the message space M to each state. The principal's strategy is a function $a_2^P : M \to [0,1]$ which assigns an action for each message received from the agent.[6]

As all the cheap talk models, there is always a babbling equilibrium in which the messages from different types of the agent are pooling together and the principal just ignores them. However, to compare to delegation, it is reasonable to focus on the informative equilibria of the communication game.

First, note that given the specific preference structure of the principal, i.e., quadratic-loss utility function, the action that the principal will take is equal to the expectation of the state conditional on the information she obtains from the agent. Since the state is binary, the principal's best action choice is simply the probability of state 1 conditional on the message sent by the agent. Therefore, the equilibrium action taken by the principal will lie in the interval $[0,1]$. In particular, in the babbling equilibrium the principal will neither update her prior about the agent's type nor about the state,

[4] Alternatively, we can think of it as the case where the babbling equilibrium is played when the principal makes a decision after communicating with the agent.

[5] Note that $\mathbb{E}[-(1/2-\theta_2)^2] = -(1/2)^2 - \mathbb{E}[\theta_2^2] + \mathbb{E}\theta_2 = -1/4$.

[6] I normalize the action choice to $[0,1]$ since I have the applications, e.g., the decision about the amount of investment given a fixed budget, in mind.

so she will take action 1/2.

Informative Equilibrium

In an informative equilibrium, the messages from the agent will induce different actions of the principal. Suppose $\bar{m}$ is the message that induces the highest action, i.e., $\bar{m} \in \text{argmax}_{m\in M} \ \Pr(\theta_2 = 1|m)$. Since the bad agent always wants the principal to choose a higher action independent of the state, he strictly prefers to send message $\bar{m}$ in each state. Formally, this strategy of the bad agent can be written as $\sigma_B(0) = \sigma_B(1) = \bar{m}$.

On the other hand, suppose $\underline{m}$ is the message that induces the lowest action, i.e., $\underline{m} \in \text{argmin}_{m\in M} \ \Pr(\theta_2 = 1|m)$. Since the good agent has the same utility function with the principal, his most preferred action is 0 when the state is 0 and 1 when the state is 1. Hence, the good agent strictly prefers to send message $\underline{m}$ when the state is 0 and message $\bar{m}$ when the state is 1. Formally, this strategy of the good agent can be written as $\sigma_G(0) = \underline{m}$ and $\sigma_G(1) = \bar{m}$.

Given the agent's strategy, upon receiving $\bar{m}$, the principal cannot infer the agent's type from the message. The principal's belief that the agent is good when receiving $\bar{m}$ is

$$\Pr(A = G|\bar{m}) = \frac{\pi_2 \Pr(\theta_2 = 1)}{\pi_2 \Pr(\theta_2 = 1) + (1 - \pi_2)} = \frac{\pi_2}{2 - \pi_2}.$$

And the principal will take action

$$a_2^P(\bar{m}) = \Pr(\theta_2 = 1|\bar{m}) = \frac{\pi_2 \Pr(\theta_2 = 1) + (1 - \pi_2)\Pr(\theta_2 = 1)}{\pi_2 \Pr(\theta_2 = 1) + (1 - \pi_2)} = \frac{1}{2 - \pi_2}.$$

In contrast, upon receiving $\underline{m}$, the principal knows that the agent must be good. That is, the principal has a belief that

$$\Pr(A = G|\underline{m}) = 1.$$

And the principal will take action

$$a_2^P(\underline{m}) = \Pr(\theta_2 = 1|\underline{m}) = 0.$$

To complete the description of the equilibrium, I assume that off the equilibrium path (i.e., upon receiving an unexpected message), the principal believes that the agent is bad and the state is 0 for sure, which leads to action 0.

Clearly, this is the unique informative equilibrium. I summarize it in the following proposition.

Proposition 3.1. *In the second period, if the principal keeps control, there exists an unique informative equilibrium in which* $\sigma_G(0) = \underline{m}$, $\sigma_G(1) = \bar{m}$, $\sigma_B(0) = \sigma_B(1) = \bar{m}$, *and* $a_2^P(\underline{m}) = 0$, $a_2^P(\bar{m}) = 1/(2 - \pi_2)$.

The equilibrium stage-payoff for both types of agents are, respectively,

$$\begin{aligned}\Gamma_G(\pi_2) =&\Pr(\theta_2 = 0)[-(a_2^P(\underline{m}) - 0)^2] + \Pr(\theta_2 = 1)[-(a_2^P(\bar{m}) - 1)^2]\\ =&-\frac{1}{2}(\frac{1-\pi_2}{2-\pi_2})^2 \end{aligned} \tag{3}$$

and

$$\begin{aligned}\Gamma_B(\pi_2) =&\Pr(\theta_2 = 0)a_2^P(\bar{m}) + \Pr(\theta_2 = 1)a_2^P(\bar{m})\\ =&\frac{1}{2-\pi_2}. \end{aligned} \tag{4}$$

Clearly, both Γ_G and Γ_B are increasing in π_2 so that the agent is incentivized to build up good reputation in the first period.

For the principal, her equilibrium stage-payoff is

$$\begin{aligned}\Gamma_P^{II}(\pi_2) =&\Pr(m = \underline{m})\Pr(\theta_2 = 0|\underline{m})[-(a_2^P(\underline{m}) - 0)^2]\\ &+ \Pr(m = \underline{m})\Pr(\theta_2 = 1|\underline{m})[-(a_2^P(\underline{m}) - 1)^2]\\ &+ \Pr(m = \bar{m})\Pr(\theta_2 = 0|\bar{m})[-(a_2^P(\bar{m}) - 0)^2]\\ &+ \Pr(m = \bar{m})\Pr(\theta_2 = 1|\bar{m})[-(a_2^P(\bar{m}) - 1)^2]\\ =&\frac{1}{2}\pi_2 \cdot 1 \cdot [-(0-0)^2] + \frac{1}{2}\pi_2 \cdot 0 \cdot [-(0-1)^2]\\ &+ (\frac{1}{2}\pi_2 + 1 - \pi_2)(1 - \frac{1}{2-\pi_2})[-(\frac{1}{2-\pi_2} - 0)^2]\\ &+ (\frac{1}{2}\pi_2 + 1 - \pi_2)\frac{1}{2-\pi_2}[-(\frac{1}{2-\pi_2} - 1)^2]\\ =&-\frac{1-\pi_2}{2(2-\pi_2)}. \end{aligned} \tag{5}$$

I will assume this unique informative equilibrium will be played in the second period when I consider the first period game. Otherwise, the agent will not care about his reputation built up in the first period if the principal's second period decision is independent of the agent's messages.

Comparative Statics

Before analyzing the principal's authority decision, it is worth pointing out some interesting comparative statics results about the second period game.

First, note that although the agent can be very biased, there always exists an informative equilibrium in which the principal's expected utility is only determined by the prior belief about the agent's type at the beginning of the second period (see (5)). That is, as long as there is some uncertainty about the agent's preferences, an informative equilibrium is guaranteed even if the agent is very misaligned with the principal. The intuition is that although the bad agent is very biased and prefers extreme actions, the good agent always strongly prefers to reveal the true state and can separate himself from the bad agent when the state is low. Given the binary structure of the state (i.e., the states are sufficiently far apart), the bad agent has no incentive to mimic the good agent's report when seeing a low state. This is in stark contrast to Crawford and Sobel (1982) who show that if the agent is too biased then no information can be communicated in equilibrium.

Second, suppose at the beginning of the second period the principal believes that the agent is very likely to be good, that is, π_2 is very close to 1. Recall that the principal would like to delegate the control only if the principal is sufficiently optimistic about the alignment of the agent (i.e., $\pi_2 \geq 1/2$, see (2)). Thus, even with an extremely biased bad agent, the principal would not like to take an uninformed decision herself. Intuitively, although the bad agent may take some very extreme actions, the proportion of the bad agent is very small (close to 0), so the principal still prefers to delegate to the agent who is very likely to have the same preferences. Moreover, at the limit, delegation and communication are doing equally well for the principal if π_2 approaches 1—the principal's expected payoff converges to 0 in both cases (see (1) and (5)).

3.3.3 Authority Allocation

We already see that if the principal communicates with the agent, there is an unique informative equilibrium. Not as other cheap talk games, I do not suffer from a plethora of equilibria. This significantly facilitates the comparison with delegation. Indeed, the setting here is simpler in an important respect than the standard model of Crawford and Sobel (1982): the state space is binary in my model while it is continuous in Crawford and Sobel (1982).

Now I am ready to determine whether the principal delegates the authority of decision making to the agent or not in the second period of the game.

Proposition 3.2. *In the second period of the game, the principal will keep the control and take a decision after communicating with the agent.*

Proof. The principal obtains expected payoff $-\frac{1-\pi_2}{2}$ from delegation while her expected payoff is $-\frac{1-\pi_2}{2(2-\pi_2)}$ in the communication equilibrium. We need to show $-\frac{1-\pi_2}{2} \leq -\frac{1-\pi_2}{2(2-\pi_2)}$. Note that if $\pi_2 = 1$, it trivially holds. To see it still holds if $\pi_2 \in [0, 1)$, note that

$$
\begin{aligned}
& -\frac{1-\pi_2}{2} < -\frac{1-\pi_2}{2(2-\pi_2)} \\
\Leftrightarrow \quad & \frac{1}{2-\pi_2} < 1
\end{aligned}
$$

which is true since $\pi_2 < 1$.

□

This proposition claims that communication is better than delegation when the principal is uncertain about the agent's motives. Assuming no uncertainty over the agent's utility function, Dessein (2002) shows that delegation is better than communication from the principal's point of view whenever informative equilibrium in the communication game is available. In contrast, the above proposition gives the opposite result, i.e., communication is preferred over delegation whenever the principal is uncertain about the agent's preferences. This coincides with Ottaviani (2000) who shows that if the agent's bias is symmetrically distributed around 0 (so the agent is unbiased in expectation), communication can improve upon delegation. However, I further show that the same result still holds even if the agent is asymmetrically biased in one direction.

In a two-action and two-state model, Garfagnini et al. (2014) show that cheap talk communication is equivalent to delegation when the agent's bias is small and the principal is uninformed about the state. Whereas, if the principal also receives a signal about the state, cheap talk communication is preferred over delegation. They arrive at these results based on the assumption that the agent's bias is common knowledge. Agastya et al. (2014) point out that if the agent is only informed of one dimension of the state, delegation does not necessarily dominate cheap talk communication or vice versa even when the agent's bias is very small.

Rush et al. (2010) also consider a framework in which the bias of the agent can only take on two possible values (which are generically not zero), the message space is assumed to consist of two elements, and the players' loss utility functions are linear rather than quadratic. They find that delegation is preferred over communication if the two types of agent are both biased upwards or downwards. Otherwise, cheap talk communication can dominate delegation. In contrast, I obtain an unambiguous result that communication is better for the principal than delegation in a different setting. In my model, one type of the agent is perfectly aligned with the principal and the message space is not restricted. Instead, the two-message equilibria are arising endogenously given the binary state.[7]

3.4 First Period of the Game

Now I go back to the first period. The good agent's payoff structure is given by

$$-(a_1 - \theta_1)^2 + \Gamma_G,$$

[7]To illustrate this, I will show in my setting for any cheap talk informative equilibrium with more than two messages, there exists a payoff-equivalent informative equilibrium in which only two messages are sent. Suppose in an informative equilibrium the set of possible messages for the agent is $\{m_1, m_2, \ldots, m_n\}$ with $n \geq 2$. Conditional on receiving message m_k, the principal will take action $\Pr(\theta_2 = 1|m_k)$ which is between 0 and 1. Without loss of generality, we can order the messages such that $\Pr(\theta_2 = 1|m_k) \leq \Pr(\theta_2 = 1|m_{k+1})$ for $k = 1, \ldots, n-1$ with at least one strict inequality. Observe that the good agent's bliss point is 0 and 1 when the state is 0 and 1, respectively. So the good agent would prefer the lowest action when the state is 0 and the highest action when the state is 1. Therefore, when the state is 0, he would be indifferent among any message that induces the lowest action. Similarly, he would be indifferent among any message that induces the highest action when the state is 1. For the bad agent, he is indifferent among any message that induces the highest action independent of the state. Clearly, we can construct another equilibrium such that the agent only sends m_1 and m_n. These equilibria are obviously payoff-equivalent.

and the bad agent's is

$$a_1 + \Gamma_B.$$

3.4.1 Delegation

If the principal delegates the decision making to the agent, the agent's action not only affects his reputation in the second period but also determines his current utility directly (In contrast, in the communication game the agent's message only indirectly affects his and the principal's current utility by influencing the principal's belief about the state of the world).

Recall that the bad agent always prefers a higher action no matter the state. The good agent may be able to take some sufficiently low action to separate from the bad agent now. Since the agent is taking action himself and only pure strategy equilibria will be considered, with a bit abuse of notation, I use $\sigma_G(\theta_1)$ and $\sigma_B(\theta_1)$ to denote the action taken by the good and bad agent for a given θ_1, respectively. Moreover, let $\pi_2(a_1)$ be the principal's belief of the agent being good for a given action a_1 of the agent.

Throughout, I will assume off the equilibrium path the principal believes the agent is surely bad. Since the agent's expected payoff of the second period is minimized at $\pi_2 = 0$ (monotonicity is not needed), this assumption is without loss of generality in the sense that given any equilibrium with some off-path beliefs we can construct another equilibrium such that the principal has the most pessimistic off-path beliefs.

To characterize the equilibrium, the following lemmas are useful.

Lemma 3.1. In any pure strategy equilibrium, if $\pi_2(\sigma_B(0)) = 0$ then $\sigma_B(0) = 1$, and if $\pi_2(\sigma_B(1)) = 0$ then $\sigma_B(1) = 1$.

Proof. First, suppose there exists some equilibrium in which $\sigma_B(0) \in [0, 1)$ and $\pi_2(\sigma_B(0)) = 0$. Assume the principal has the most pessimistic belief off the equilibrium path. By taking action 1 in state 0, the bad agent is strictly better off in terms of current utility and he obtains a weakly higher reputation than taking action $\sigma_B(0)$, i.e., $\pi_2(1) \geq \pi_2(\sigma_B(0)) = 0$. If the principal is less pessimistic when observing action 1, the bad agent is strictly better off even in terms of reputation. Hence, we have found a profitable deviation for the bad agent when the state is 0, which unravels the equilibrium.

By the same logic, it can be shown that if $\pi_2(\sigma_B(1)) = 0$ then $\sigma_B(1) = 1$.

□

Lemma 3.2. In any pure strategy equilibrium, if $\pi_2(\sigma_G(0)) \leq \pi_2(\sigma_G(1))$, then $\sigma_G(0) \leq \sigma_G(1)$.

Proof. I show the contrapositive. Suppose $1 \geq \sigma_G(0) > \sigma_G(1) \geq 0$ (note that the action space is $[0, 1]$). If $\pi_2(\sigma_G(0)) \leq \pi_2(\sigma_G(1))$, then the good agent will deviate to action $\sigma_G(1)$ when the state θ_1 is 0 since the action $\sigma_G(1)$ yields the good agent weakly higher reputation and strictly higher current utility (when the state is 0 the good agent's bliss point is 0), which is a contradiction. □

Lemma 3.3. In any pure strategy equilibrium, $\sigma_G(0) \neq \sigma_G(1)$.

Proof. Suppose by contradiction that $\sigma_G(0) = \sigma_G(1)$ in some pure strategy equilibrium. There are four cases to be considered.

(i) $\pi_2(\sigma_B(0)) = \pi_2(\sigma_B(1)) = 0$. By Lemma 3.1, it implies that this case corresponds to the equilibrium where $\sigma_B(0) = \sigma_B(1) = 1$, $0 \leq \sigma_G(0) = \sigma_G(1) < 1$, and $\pi_2(\sigma_G(0)) = \pi_2(\sigma_G(1)) = 1$. Then the following incentive constraints need to be satisfied:

$$
\begin{aligned}
-(\sigma_G(0) - 0)^2 + \Gamma_G(\pi_2(\sigma_G(0))) &\geq -(0 - 0)^2 + \Gamma_G(0) \\
-(\sigma_G(1) - 1)^2 + \Gamma_G(\pi_2(\sigma_G(1))) &\geq -(1 - 1)^2 + \Gamma_G(0)
\end{aligned}
$$

The first inequality requires that the good agent has no incentive to deviate to action 0 when the state is 0 which induces the principal to believe the agent is surely bad, i.e., $\pi_2(0) = 0$ (we still assume the most pessimistic belief off the equilibrium path). If $\sigma_G(0)$ happens to be 0, the inequality still holds since $\Gamma_G(\pi_2(\sigma_G(0))) = \Gamma_G(1) > \Gamma_G(0)$. Similarly, the second inequality requires that the good agent has no incentive to deviate to action 1 when the state is 1. However, by simple algebra there is no $\sigma_G(0) = \sigma_G(1)$ such that these two constraints hold simultaneously, which is a contradiction to the assumption that $\sigma_G(0)$ and $\sigma_G(1)$ are the good agent's equilibrium strategy. To be clear about this, note that $\Gamma_G(1) = 0$ and $\Gamma_G(0) = -\frac{1}{8}$. Then the two inequalities can be rewritten as

$$-\frac{1}{2\sqrt{2}} \leq \sigma_G(0) \leq \frac{1}{2\sqrt{2}}$$
$$1-\frac{1}{2\sqrt{2}} \leq \sigma_G(0) \leq 1+\frac{1}{2\sqrt{2}}$$

but $\frac{1}{2\sqrt{2}} < 1-\frac{1}{2\sqrt{2}}$ and therefore there is no intersection between $[-\frac{1}{2\sqrt{2}}, \frac{1}{2\sqrt{2}}]$ and $[1-\frac{1}{2\sqrt{2}}, 1+\frac{1}{2\sqrt{2}}]$.

(ii) $\pi_2(\sigma_B(0)) > 0, \pi_2(\sigma_B(1)) > 0$. This case corresponds to the pooling equilibrium where $\sigma_G(0) = \sigma_G(1) = \sigma_B(0) = \sigma_B(1)$ and $\pi_2(\sigma_G(0)) = \pi_2(\sigma_G(1)) = \pi_1$. If this strategy profile constitutes an equilibrium, then the following incentive constraints for the good agent need to be satisfied:

$$-(\sigma_G(0)-0)^2 + \Gamma_G(\pi_2(\sigma_G(0))) \geq -(0-0)^2 + \Gamma_G(0)$$
$$-(\sigma_G(1)-1)^2 + \Gamma_G(\pi_2(\sigma_G(1))) \geq -(1-1)^2 + \Gamma_G(0)$$

where $\Gamma_G(\pi_2(\sigma_G(0))) = \Gamma_G(\pi_2(\sigma_G(1))) = \Gamma_G(\pi_1) = -\frac{1}{2}(\frac{1-\pi_1}{2-\pi_1})^2$ and $\Gamma_G(0) = -\frac{1}{8}$. Plugging in $\Gamma_G(\cdot)$, we have

$$-\sqrt{\frac{1}{8}-\frac{1}{2}(\frac{1-\pi_1}{2-\pi_1})^2} \leq \sigma_G(0) \leq \sqrt{\frac{1}{8}-\frac{1}{2}(\frac{1-\pi_1}{2-\pi_1})^2}$$
$$1-\sqrt{\frac{1}{8}-\frac{1}{2}(\frac{1-\pi_1}{2-\pi_1})^2} \leq \sigma_G(0) \leq 1+\sqrt{\frac{1}{8}-\frac{1}{2}(\frac{1-\pi_1}{2-\pi_1})^2}$$

Since $\sqrt{\frac{1}{8}-\frac{1}{2}(\frac{1-\pi_1}{2-\pi_1})^2} \leq \sqrt{\frac{1}{8}} < \frac{1}{2}$, we have $1-\sqrt{\frac{1}{8}-\frac{1}{2}(\frac{1-\pi_1}{2-\pi_1})^2} > \sqrt{\frac{1}{8}-\frac{1}{2}(\frac{1-\pi_1}{2-\pi_1})^2}$ so that there is no $\sigma_G(0)$ such that the incentive constraints for the good agent hold simultaneously, which is a contradiction.

(iii) $\pi_2(\sigma_B(0)) > 0, \pi_2(\sigma_B(1)) = 0$. This case corresponds to the equilibrium where $0 \leq \sigma_G(0) = \sigma_G(1) = \sigma_B(0) < 1, \sigma_B(1) = 1$, and $\pi_2(\sigma_G(0)) = \pi_2(\sigma_G(1)) = \frac{2\pi_1}{1+\pi_1}$. If this strategy profile constitutes an equilibrium, then the following incentive constraints for the good agent need to be satisfied:

$$-(\sigma_G(0)-0)^2 + \Gamma_G(\pi_2(\sigma_G(0))) \geq -(0-0)^2 + \Gamma_G(0)$$
$$-(\sigma_G(1)-1)^2 + \Gamma_G(\pi_2(\sigma_G(1))) \geq -(1-1)^2 + \Gamma_G(0)$$

where $\Gamma_G(\pi_2(\sigma_G(0))) = \Gamma_G(\pi_2(\sigma_G(1))) = \Gamma_G(\frac{2\pi_1}{1+\pi_1}) = -\frac{1}{8}(1-\pi_1)^2$ and $\Gamma_G(0) = -\frac{1}{8}$. Plugging in $\Gamma_G(\cdot)$, we have

$$-\sqrt{\frac{1}{8} - \frac{1}{8}(1-\pi_1)^2} \leq \sigma_G(0) \leq \sqrt{\frac{1}{8} - \frac{1}{8}(1-\pi_1)^2}$$
$$1 - \sqrt{\frac{1}{8} - \frac{1}{8}(1-\pi_1)^2} \leq \sigma_G(0) \leq 1 + \sqrt{\frac{1}{8} - \frac{1}{8}(1-\pi_1)^2}$$

Since $\sqrt{\frac{1}{8} - \frac{1}{8}(1-\pi_1)^2} \leq \sqrt{\frac{1}{8}} < \frac{1}{2}$, we have $1 - \sqrt{\frac{1}{8} - \frac{1}{8}(1-\pi_1)^2} > \sqrt{\frac{1}{8} - \frac{1}{8}(1-\pi_1)^2}$ so that there is no $\sigma_G(0)$ such that the incentive constraints for the good agent hold simultaneously, which is a contradiction.

(iv) $\pi_2(\sigma_B(0)) = 0, \pi_2(\sigma_B(1)) > 0$. This case corresponds to the equilibrium where $0 \leq \sigma_G(0) = \sigma_G(1) = \sigma_B(1) < 1, \sigma_B(0) = 1$ and $\pi_2(\sigma_G(0)) = \pi_2(\sigma_G(1)) = \frac{2\pi_1}{1+\pi_1}$. The argument is similar to case (iii), which is omitted. □

This lemma implies that there is no pooling equilibrium.

Lemma 3.4. In any pure strategy equilibrium, if $\pi_2(\sigma_G(0)) \leq \pi_2(\sigma_G(1))$, then $\sigma_G(0) \neq \sigma_B(1)$.

Proof. Suppose by contradiction that $\sigma_G(0) = \sigma_B(1)$ in some equilibrium. This implies that $\pi_2(\sigma_B(1)) > 0$. Hence, we need to consider two cases: (i) $\pi_2(\sigma_B(0)) = 0$ and (ii) $\pi_2(\sigma_B(0)) > 0$.

First consider the former case. By Lemma 3.1, $\pi_2(\sigma_B(0)) = 0$ implies $\sigma_B(0) = 1$. Moreover, by Lemma 3.2 and 3.3, we must have $\sigma_B(1) = \sigma_G(0) < \sigma_G(1)$, $\sigma_G(1) \in [0, 1)$, $\pi_2(\sigma_G(0)) = \pi_2(\sigma_B(1)) = \pi_1$ and $\pi_2(\sigma_G(1)) = 1$. Hence, the bad agent would like to deviate to action $\sigma_G(1)$ when the state is 1 since it yields a higher current utility ($\sigma_B(1) < \sigma_G(1)$) and also a higher reputation because $\pi_2(\sigma_B(1)) = \pi_1 < 1 = \pi_2(\sigma_G(1))$ (recall that $\Gamma_B(\pi_2)$ is increasing in π_2), a contradiction.

For the latter case, we need to distinguish two subcases: (1) $\sigma_G(0) = \sigma_B(1) = \sigma_B(0)$ and (2) $\sigma_G(0) = \sigma_B(1) \neq \sigma_B(0)$. Consider the first subcase. By Lemma 3.2 and 3.3, we know that $\sigma_G(0) = \sigma_B(0) = \sigma_B(1) < \sigma_G(1)$. Therefore, by Bayes' rule $\pi_2(\sigma_B(1)) = \frac{\pi_1}{2-\pi_1} < 1 = \pi_2(\sigma_G(1))$. By the same argument as the former case, the bad agent will be strictly better off by deviating to action $\sigma_G(1)$ no matter the state, a contradiction.

For the second subcase, $\pi_2(\sigma_B(0)) > 0$ implies that $\sigma_B(0) = \sigma_G(1)$. Hence, by Bayes' rule $\pi_2(\sigma_G(0)) = \pi_2(\sigma_B(1)) = \pi_2(\sigma_B(0)) = \pi_2(\sigma_G(1)) = \pi_1$. By Lemma 3.2 and

3.3, we must have $\sigma_G(0) < \sigma_G(1)$ and so $\sigma_B(1) < \sigma_G(1)$. Hence, the bad agent would like to deviate to action $\sigma_G(1)$ when the state is 1 since it yields the same reputation ($\pi_2(\sigma_B(1)) = \pi_2(\sigma_G(1)) = \pi_1$) but a higher current utility because $\sigma_B(1) < \sigma_G(1)$, a contradiction. □

Lemma 3.5. In any pure strategy equilibrium, if $\pi_2(\sigma_G(0)) \leq \pi_2(\sigma_G(1))$, then $\sigma_G(0) \neq \sigma_B(0)$.

Proof. Suppose by contradiction that $\sigma_G(0) = \sigma_B(0)$ in some equilibrium. This implies that $\pi_2(\sigma_B(0)) > 0$. Hence, we need to consider two cases: (i) $\pi_2(\sigma_B(1)) = 0$ and (ii) $\pi_2(\sigma_B(1)) > 0$.

First consider the former case. By Lemma 3.1, $\pi_2(\sigma_B(1)) = 0$ implies $\sigma_B(1) = 1$. Moreover, by Lemma 3.2 and 3.3, we must have $\sigma_B(0) = \sigma_G(0) < \sigma_G(1)$, $\sigma_G(1) \in [0, 1)$, $\pi_2(\sigma_G(0)) = \pi_2(\sigma_B(0)) = \pi_1$ and $\pi_2(\sigma_G(1)) = 1$. Hence, the bad agent would like to deviate to action $\sigma_G(1)$ when the state is 0 since it yields a higher current utility ($\sigma_B(0) < \sigma_G(1)$) and also a higher reputation because $\pi_2(\sigma_B(0)) = \pi_1 < 1 = \pi_2(\sigma_G(1))$ (recall that $\Gamma_B(\pi_2)$ is increasing in π_2), a contradiction.

For the latter case, by Lemma 3.4, $\sigma_G(0) \neq \sigma_B(1)$. Thus, $\pi_2(\sigma_B(1)) > 0$ implies that $\sigma_B(1) = \sigma_G(1)$. By Lemma 3.2 and 3.3, we must have $\sigma_G(0) < \sigma_G(1)$ and so $\sigma_B(0) < \sigma_B(1)$. Hence, by Bayes' rule $\pi_2(\sigma_G(0)) = \pi_2(\sigma_G(1)) = \pi_2(\sigma_B(0)) = \pi_2(\sigma_B(1)) = \pi_1$. Therefore, the bad agent would like to deviate to action $\sigma_B(1)$ when the state is 0 since it yields the same reputation ($\pi_2(\sigma_B(0)) = \pi_2(\sigma_B(1)) = \pi_1$) but a higher current utility because $\sigma_B(0) < \sigma_B(1)$, a contradiction.

□

Lemma 3.6. In any pure strategy equilibrium, $\pi_2(\sigma_B(0)) > 0$ or $\pi_2(\sigma_B(1)) > 0$.

Proof. Suppose by contradiction that $\pi_2(\sigma_B(0)) = \pi_2(\sigma_B(1)) = 0$ in some equilibrium. It implies that $\sigma_B(0) = \sigma_B(1) = 1$ and $\pi_2(\sigma_G(0)) = \pi_2(\sigma_G(1)) = 1$. By Lemma 3.2 and 3.3, we must have $0 \leq \sigma_G(0) < \sigma_G(1) < 1$.

If there is such an equilibrium, the following incentive constraints need to be satisfied:

$$\sigma_B(1) + \Gamma_B(\pi_2(\sigma_B(1))) \geq \sigma_G(1) + \Gamma_B(\pi_2(\sigma_G(1)))$$
$$-(\sigma_G(1) - 1)^2 + \Gamma_G(\pi_2(\sigma_G(1))) \geq -(\sigma_B(1) - 1)^2 + \Gamma_G(\pi_2(\sigma_B(1)))$$

The first inequality requires that the bad agent has no incentive to deviate to action $\sigma_G(1)$ when the state is 1 and the second inequality requires that the good agent has no incentive to deviate to action $\sigma_B(1)$ when the state is 1. Plugging in $\sigma_B(1) = 1$ and $\pi_2(\cdot)$ and rearranging, we have

$$\sigma_G(1) \leq \frac{1}{2} \text{ and } 1 - \frac{1}{2\sqrt{2}} \leq \sigma_G(1) < 1.$$

However, since $\frac{1}{2} < 1 - \frac{1}{2\sqrt{2}}$ the two incentive constraints cannot hold simultaneously, which is a contradiction. □

Combining all the lemmas above, the equilibrium candidates can be narrowed down to three cases which are summarized in the following proposition.

Proposition 3.3. *In the first period, if the principal delegates the decision making to the agent, a pure strategy equilibrium must feature one of the following structures:*

(i) $\sigma_G(0) < \sigma_G(1) = \sigma_B(0) < \sigma_B(1) = 1$*;*

(ii) $\sigma_G(0) < \sigma_G(1) = \sigma_B(1) < \sigma_B(0) = 1$*;*

(iii) $\sigma_G(0) < \sigma_G(1) = \sigma_B(0) = \sigma_B(1)$.

Proof. By Lemma 3.6, in any pure strategy equilibrium, we must have $\pi_2(\sigma_B(0)) > 0$ or $\pi_2(\sigma_B(1)) > 0$. First, if $\pi_2(\sigma_B(0)) > 0$ and $\pi_2(\sigma_B(1)) = 0$, then $\sigma_B(0) = \sigma_G(0)$ or $\sigma_B(0) = \sigma_G(1)$, and by Lemma 3.1 $\sigma_B(1) = 1$. Moreover, we must have $\sigma_B(0) = \sigma_G(1) > \sigma_G(0)$. To see this, note that by Lemma 3.3 $\sigma_G(0) \neq \sigma_G(1)$ and if $\sigma_B(0) = \sigma_G(0)$ then we must have $\pi_2(\sigma_G(0)) = \pi_2(\sigma_B(0)) = \pi_1$ and $\pi_2(\sigma_G(1)) = 1$ (if $\pi_2(\sigma_G(1)) < 1$, it implies that $\sigma_G(1) = \sigma_B(1)$ and therefore $\pi_2(\sigma_B(1)) > 0$, which is a contradiction). By Lemma 3.5, if $\pi_2(\sigma_G(0)) \leq \pi_2(\sigma_G(1))$, then $\sigma_G(0) \neq \sigma_B(0)$, a contradiction. To see $\sigma_G(0) < \sigma_G(1)$, note that since $\sigma_G(0) \neq \sigma_G(1)$ and $\sigma_B(0) = \sigma_G(1) < 1 = \sigma_B(1)$ we have $\pi_2(\sigma_G(0)) = 1 > \pi_1 = \pi_2(\sigma_B(0))$, so if $\sigma_G(0) > \sigma_G(1)$ the bad agent will deviate to action $\sigma_G(0)$ when the state is 0 because $\sigma_G(0)$ yields a higher current utility and also a higher reputation. This proves the first possible structure that a pure strategy equilibrium can take.

Second, if $\pi_2(\sigma_B(0)) = 0$ and $\pi_2(\sigma_B(1)) > 0$, then $\sigma_B(1) = \sigma_G(0)$ or $\sigma_B(1) = \sigma_G(1)$, and by Lemma 3.1 $\sigma_B(0) = 1$. Moreover, we must have $\sigma_B(1) = \sigma_G(1) > \sigma_G(0)$. To see this, note that by Lemma 3.3 $\sigma_G(0) \neq \sigma_G(1)$ and if $\sigma_B(1) = \sigma_G(0)$ then we must have $\pi_2(\sigma_G(0)) = \pi_2(\sigma_B(1)) = \pi_1$ and $\pi_2(\sigma_G(1)) = 1$ (if $\pi_2(\sigma_G(1)) < 1$, it implies that $\sigma_G(1) = \sigma_B(0)$ and therefore $\pi_2(\sigma_B(0)) > 0$, which is a contradiction). By

Lemma 3.4, if $\pi_2(\sigma_G(0)) \leq \pi_2(\sigma_G(1))$, then $\sigma_G(0) \neq \sigma_B(1)$, a contradiction. To see $\sigma_G(0) < \sigma_G(1)$, note that since $\sigma_G(0) \neq \sigma_G(1)$ and $\sigma_B(1) = \sigma_G(1) < 1 = \sigma_B(0)$ we have $\pi_2(\sigma_G(0)) = 1 > \pi_1 = \pi_2(\sigma_B(1))$, so if $\sigma_G(0) > \sigma_G(1)$ the bad agent will deviate to action $\sigma_G(0)$ when the state is 1 because $\sigma_G(0)$ yields a higher current utility and also a higher reputation. This proves the second possible structure that a pure strategy equilibrium can take.

Third, if $\pi_2(\sigma_B(0)) > 0$ and $\pi_2(\sigma_B(1)) > 0$, we must have $\sigma_G(0) \neq \sigma_B(1)$. To see this, note that if $\sigma_G(0) = \sigma_B(1)$, by Lemma 3.3 $\sigma_G(0) \neq \sigma_G(1)$, so we must have $\pi_2(\sigma_G(0)) = \frac{\pi_1}{2-\pi_1}$ and $\pi_2(\sigma_G(1)) = 1$ with $\sigma_B(0) = \sigma_G(0)$, or $\pi_2(\sigma_G(0)) = \pi_2(\sigma_G(1)) = \pi_1$ with $\sigma_B(0) = \sigma_G(1)$. In either case, $\pi_2(\sigma_G(0)) \leq \pi_2(\sigma_G(1))$ holds, which implies $\sigma_G(0) \neq \sigma_B(1)$ by Lemma 3.4, which is a contradiction. Moreover, we must have $\sigma_B(0) = \sigma_B(1)$. If not, $\pi_2(\sigma_B(1)) > 0$ implies $\sigma_G(1) = \sigma_B(1)$ since $\sigma_G(0) \neq \sigma_B(1)$, and therefore $\sigma_G(0) = \sigma_B(0)$ since $\pi_2(\sigma_B(0)) > 0$. Thus, we have $\pi_2(\sigma_G(0)) = \pi_2(\sigma_G(1)) = \pi_1$, which implies $\sigma_G(0) \neq \sigma_B(0)$ by Lemma 3.5, which is a contradiction. Hence, we have $\sigma_B(0) = \sigma_B(1) = \sigma_G(1)$ since $\pi_2(\sigma_B(0)) > 0$ and $\pi_2(\sigma_B(1)) > 0$. Last, to see $\sigma_G(0) < \sigma_G(1)$, note that since $\sigma_G(0) \neq \sigma_G(1)$ and $\sigma_B(0) = \sigma_B(1) = \sigma_G(1)$, we have $\pi_2(\sigma_G(0)) = 1 > \frac{\pi_1}{2-\pi_1} = \pi_2(\sigma_G(1))$, so if $\sigma_G(0) > \sigma_G(1)$ the bad agent will deviate to action $\sigma_G(0)$ no matter the state because $\sigma_G(0)$ yields a higher current utility and also a higher reputation. This proves the third possible structure that a pure strategy equilibrium can take. □

This proposition shows that if the agent has the authority to make decisions, the bad agent can behave sensitively to the state as the good agent does, which provides possibilities of delegation improving upon communication when the agent has reputational concerns. However, it is worth noting that this reputation effect is still not strong enough to refrain the bad agent from taking high actions. In particular, it fails to discipline the bad agent to take the principal's desired action in each state so that the principal cannot achieve the first-best outcome (i.e., $(\sigma_G(0), \sigma_G(1), \sigma_B(0), \sigma_B(1)) = (0, 1, 0, 1)$ cannot be an equilibrium).

3.4.2 Communication

As argued before, we can without loss of generality focus on the case where the message space consists of two elements, $M = \{\underline{m}, \bar{m}\}$. In an informative equilibrium, assume $\bar{m}$ induces a higher action. Likewise, the action that the principal will take when receiving message m is equal to the conditional probability she assigns to state 1. Therefore, we essentially assume $\Pr(\theta_1 = 1|\bar{m}) > \Pr(\theta_1 = 1|\underline{m})$.

Note that the expected payoff of the agent (whatever his type) in the second period is a function of π_2, namely, the reputation of the agent at the beginning of the second period, which in turn depends on how the principal updates her belief upon receiving the message from the agent in the first period.

To characterize the equilibrium, the following lemmas are useful.

Lemma 3.7. Let $\pi_2(m)$ denote the probability of the agent being good when the principal receives message m. Then in any informative equilibrium, we have $\pi_2(\underline{m}) \geq \pi_2(\bar{m})$.

Proof. I prove by contradiction. Suppose $\pi_2(\underline{m}) < \pi_2(\bar{m})$ in some equilibrium. Recall that the bad agent always prefers a higher action no matter the state. Since $\Pr(\theta_1 = 1|\bar{m}) > \Pr(\theta_1 = 1|\underline{m})$ by assumption, the bad agent derives higher current utility from reporting $\bar{m}$. Moreover, the bad agent's expected payoff of the second period $\Gamma_B(\pi_2)$ is increasing in π_2. Hence, it is strictly better for the bad agent to send message $\bar{m}$. Namely, we have $\sigma_B(0) = \sigma_B(1) = \bar{m}$.

The principal updates her belief by Bayes' rule whenever possible. If the probability is not well defined (i.e., the denominator is zero), the convention that the principal retains her prior belief π_1 is adopted.

Now if $\sigma_G(0) = \sigma_G(1) = \bar{m}$, then

$$\pi_2(\bar{m}) = \pi_1 = \pi_2(\underline{m})$$

which is a contradiction.

If $\sigma_G(0) \neq \bar{m}$ or $\sigma_G(1) \neq \bar{m}$, then

$$\pi_2(\underline{m}) = 1$$

which implies $\pi_2(\bar{m}) > 1$, a contradiction. □

Lemma 3.8. In any informative equilibrium, the good agent will report $\underline{m}$ when the state is 0. That is, $\sigma_G(0) = \underline{m}$.

Proof. By Lemma 3.7, we know that the good agent will obtain a higher reputation by sending message $\underline{m}$. Note that the expected payoff of the good agent in the second period $\Gamma_G(\pi_2)$ is increasing in π_2. Moreover, the good agent's bliss point is 0 when the state is 0, so given that $\Pr(\theta_1 = 1|\bar{m}) > \Pr(\theta_1 = 1|\underline{m})$, the good agent will obtain a higher current utility from sending message $\underline{m}$. Overall, it is strictly better for the good agent to send message $\underline{m}$. Namely, $\sigma_G(0) = \underline{m}$.

□

Lemma 3.9. In any pure strategy informative equilibrium, the good agent will report $\bar{m}$ when the state is 1. That is, $\sigma_G(1) = \bar{m}$.

Proof. By Lemma 3.8, we know that $\sigma_G(0) = \underline{m}$ in any informative equilibrium. I will show it cannot be an equilibrium that (i) $(\sigma_G(0), \sigma_G(1), \sigma_B(0), \sigma_B(1)) = (\underline{m}, \underline{m}, \bar{m}, \underline{m})$, (ii) $(\sigma_G(0), \sigma_G(1), \sigma_B(0), \sigma_B(1)) = (\underline{m}, \underline{m}, \bar{m}, \bar{m})$, and (iii) $(\sigma_G(0), \sigma_G(1), \sigma_B(0), \sigma_B(1)) = (\underline{m}, \underline{m}, \underline{m}, \bar{m})$.

I start with the first case. If $(\sigma_G(0), \sigma_G(1), \sigma_B(0), \sigma_B(1)) = (\underline{m}, \underline{m}, \bar{m}, \underline{m})$ is an equilibrium, then it implies that $\Pr(\theta_1 = 1|\bar{m}) = 0$ and

$$\begin{aligned}\Pr(\theta_1 = 1|\underline{m}) &= \frac{\frac{1}{2}\pi_1 + \frac{1}{2}(1-\pi_1)}{\pi_1 + \frac{1}{2}(1-\pi_1)} = \frac{1}{1+\pi_1} \\ &> 0 = \Pr(\theta_1 = 1|\bar{m}),\end{aligned}$$

which is a contradiction to the assumption that $\Pr(\theta_1 = 1|\bar{m}) >\Pr(\theta_1 = 1|\underline{m})$.

Consider the second case. If $(\sigma_G(0), \sigma_G(1), \sigma_B(0), \sigma_B(1)) = (\underline{m}, \underline{m}, \bar{m}, \bar{m})$ is an equilibrium, it implies that

$$\Pr(\theta_1 = 1|\bar{m}) = \Pr(\theta_1 = 1|\underline{m}) = \frac{1}{2},$$

which is also a contradiction to the assumption that $\Pr(\theta_1 = 1|\bar{m}) >\Pr(\theta_1 = 1|\underline{m})$.

At last, suppose $(\sigma_G(0), \sigma_G(1), \sigma_B(0), \sigma_B(1)) = (\underline{m}, \underline{m}, \underline{m}, \bar{m})$ is an equilibrium. Then we must have

$$\Pr(\theta_1 = 1|\underline{m}) = \frac{\pi_1}{1+\pi_1} \text{ and } \Pr(\theta_1 = 1|\bar{m}) = 1.$$

That is, the principal will take action $\frac{\pi_1}{1+\pi_1}$ upon receiving message $\underline{m}$ and action 1 when receiving message $\bar{m}$. Moreover, the updated beliefs of the principal as a response to the messages are

$$\pi_2(\underline{m}) = \frac{2\pi_1}{1+\pi_1} \text{ and } \pi_2(\bar{m}) = 0.$$

Given the principal and the bad agent's strategies, the good agent should have no incentive to send message $\bar{m}$ when the state θ_1 is 1. That is, the following incentive constraint for the good agent must hold:

$$-(a_1^P(\underline{m}) - 1)^2 + \Gamma_G(\pi_2(\underline{m})) \geq -(a_1^P(\bar{m}) - 1)^2 + \Gamma_G(\pi_2(\bar{m})).$$

Hence,

$$-(\frac{\pi_1}{1+\pi_1} - 1)^2 - \frac{1}{2}(\frac{1 - \frac{2\pi_1}{1+\pi_1}}{2 - \frac{2\pi_1}{1+\pi_1}})^2 \geq -(1-1)^2 - \frac{1}{2}(\frac{1-0}{2-0})^2.$$

It implies that

$$(1+\pi_1)^2[1 - (1-\pi_1)^2] \geq 8,$$

which cannot be true since $0 \leq (1+\pi_1)^2 \leq 4$ and $0 \leq 1 - (1-\pi_1)^2 \leq 1$. □

Combining all the lemmas above, we obtain the following unique equilibrium that features the same structures as the communication equilibrium in the second period.

Proposition 3.4. *In the first period, if the principal keeps control, there exists an unique pure strategy informative equilibrium in which* $(\sigma_G(0), \sigma_G(1), \sigma_B(0), \sigma_B(1)) = (\underline{m}, \bar{m}, \bar{m}, \bar{m})$.

Proof. By Lemma 3.8 and 3.9, we know that $\sigma_G(0) = \underline{m}$ and $\sigma_G(1) = \bar{m}$ in any pure strategy informative equilibrium. I first verify that $(\sigma_G(0), \sigma_G(1), \sigma_B(0), \sigma_B(1)) = (\underline{m}, \bar{m}, \bar{m}, \bar{m})$ is an equilibrium. And then I will show it cannot be an equilibrium that (i) $(\sigma_G(0), \sigma_G(1), \sigma_B(0), \sigma_B(1)) = (\underline{m}, \bar{m}, \underline{m}, \underline{m})$, (ii) $(\sigma_G(0), \sigma_G(1), \sigma_B(0), \sigma_B(1)) = (\underline{m}, \bar{m}, \underline{m}, \bar{m})$, and (iii) $(\sigma_G(0), \sigma_G(1), \sigma_B(0), \sigma_B(1)) = (\underline{m}, \bar{m}, \bar{m}, \underline{m})$.

To verify $(\sigma_G(0), \sigma_G(1), \sigma_B(0), \sigma_B(1)) = (\underline{m}, \bar{m}, \bar{m}, \bar{m})$ is an equilibrium, note that given this strategy profile the principal's best response is to take action

$$a_1^P(\bar{m}) = \Pr(\theta_1 = 1|\bar{m}) = \frac{1}{2-\pi_1}$$

when receiving message $\bar{m}$ and action $a_1^P(\underline{m})$ =$\Pr(\theta_1 = 1|\underline{m}) = 0$ upon receiving message $\underline{m}$. Moreover, by Bayes' rule, the principal believes that the agent is good with probability

$$\pi_2(\bar{m}) = \frac{\pi_1}{2-\pi_1}$$

when receiving message $\bar{m}$ and with probability $\pi_2(\underline{m}) = 1$ upon receiving message $\underline{m}$.

Clearly, the good agent has no incentive to deviate when the state is 0. Hence, we have to check the following three incentive constraints: (1) the good agent has no incentive to send message $\underline{m}$ when the state is 1 so that

$$-(a_1^P(\bar{m})-1)^2 + \Gamma_G(\pi_2(\bar{m})) \geq -(a_1^P(\underline{m})-1)^2 + \Gamma_G(\pi_2(\underline{m})),$$

plugging in $a_1^P(\cdot)$ and $\pi_2(\cdot)$ and rearranging we have

$$(\frac{1-\pi_1}{2-\pi_1})^2 + 2(\frac{1-\pi_1}{4-3\pi_1})^2 \leq 1,$$

which is true since the function $\pi_1 \mapsto (\frac{1-\pi_1}{2-\pi_1})^2 + 2(\frac{1-\pi_1}{4-3\pi_1})^2$ is decreasing in π_1 and therefore

$$(\frac{1-\pi_1}{2-\pi_1})^2 + 2(\frac{1-\pi_1}{4-3\pi_1})^2 \leq \frac{3}{8} < 1.$$

(2) the bad agent has no incentive to send message $\underline{m}$ no matter the state so that

$$a_1^P(\bar{m}) + \Gamma_B(\pi_2(\bar{m})) \geq a_1^P(\underline{m}) + \Gamma_B(\pi_2(\underline{m})),$$

plugging in $a_1^P(\cdot)$ and $\pi_2(\cdot)$ and rearranging we have

$$\frac{1}{2-\pi_1} + \frac{2-\pi_1}{4-3\pi_1} \geq 1$$

which is true since the function $\pi_1 \mapsto \frac{1}{2-\pi_1} + \frac{2-\pi_1}{4-3\pi_1}$ is increasing in π_1.

Now I will show the uniqueness by ruling out the three cases listed above in turn. I start with the first case. If $(\sigma_G(0), \sigma_G(1), \sigma_B(0), \sigma_B(1)) = (\underline{m}, \bar{m}, \underline{m}, \underline{m})$ is an equilibrium, then it implies that $\pi_2(\bar{m}) = 1$ and

$$\pi_2(\underline{m}) = \frac{\frac{1}{2}\pi_1}{\frac{1}{2}\pi_1 + (1-\pi_1)} = \frac{\pi_1}{2-\pi_1} < 1 = \pi_2(\bar{m}),$$

which is a contradiction to Lemma 3.7.

Consider the second case. Suppose $(\sigma_G(0), \sigma_G(1), \sigma_B(0), \sigma_B(1)) = (\underline{m}, \bar{m}, \underline{m}, \bar{m})$ is an equilibrium. Given the agent's strategy, the principal retains her prior belief of the agent being good, i.e., $\pi_2(\underline{m}) = \pi_2(\bar{m}) = \pi_1$. Hence, the agent will obtain the same reputation no matter which message he announces, which implies that the agent's expected payoff of the second period does not depend on the message. However, recall that the bad agent prefers a higher decision no matter the state. Now the bad agent has an incentive to deviate to message $\bar{m}$ when the state is 0 since $\bar{m}$ is assumed to induce a higher action.

At last, if $(\sigma_G(0), \sigma_G(1), \sigma_B(0), \sigma_B(1)) = (\underline{m}, \bar{m}, \bar{m}, \underline{m})$, then we must have

$$\pi_2(\underline{m}) = \pi_2(\bar{m}) = \frac{\frac{1}{2}\pi_1}{\frac{1}{2}\pi_1 + \frac{1}{2}(1-\pi_1)} = \pi_1.$$

However, if there is no reputational cost of sending message $\bar{m}$, the bad agent will always report $\bar{m}$ no matter the state since message $\bar{m}$ induces a higher action, which is in favor of the bad agent, a contradiction. □

Given the unique pure strategy informative equilibrium, the principal's first period expected payoff in the communication game can be derived:

$$\begin{aligned}\Gamma_P^I(\pi_1) =& \Pr(m = \underline{m})\Pr(\theta_1 = 0|\underline{m})[-(a_1^P(\underline{m}) - 0)^2] \\ &+ \Pr(m = \underline{m})\Pr(\theta_1 = 1|\underline{m})[-(a_1^P(\underline{m}) - 1)^2] \\ &+ \Pr(m = \bar{m})\Pr(\theta_1 = 0|\bar{m})[-(a_1^P(\bar{m}) - 0)^2] \\ &+ \Pr(m = \bar{m})\Pr(\theta_1 = 1|\bar{m})[-(a_1^P(\bar{m}) - 1)^2] \\ =& -\frac{1-\pi_1}{2(2-\pi_1)}.\end{aligned}$$

3.4.3 Authority Allocation

In the first period, if the principal takes an action after communicating with the agent, there exists an unique pure strategy informative equilibrium. However, if the principal delegates the decision making to the agent, three types of pure strategy equilibrium can arise. In order to compare this two decision making protocols, it is sufficient to focus on the principal-optimal equilibrium out of all the equilibria for each parameter configuration.

Proposition 3.5. *In the first period, there exists a prior cut-off* $\bar{\pi}_1$ *such that communication dominates delegation if* $\pi_1 \leq \bar{\pi}_1$ *and delegation dominates communication if* $\pi_1 > \bar{\pi}_1$.

Proof. In the first period, if the principal keeps control, her expected payoff from the unique pure strategy informative equilibrium is

$$\Gamma_P^I(\pi_1) = -\frac{1-\pi_1}{2(2-\pi_1)}.$$

And recall that the principal will choose to communicate with the agent in the second period, which yields the principal expected payoff

$$\Gamma_P^{II}(\pi_2) = -\frac{1-\pi_2}{2(2-\pi_2)}.$$

Since the agent's reputation π_2 is determined by the first period's play, which will affect the principal's expected payoff of the second period, we first calculate this payoff from a given equilibrium in the first period. Specifically, if type (i) or (ii) delegation equilibrium is played in the first period, the principal's expected payoff of the second period is

$$\begin{aligned}\tilde{\Gamma}_P^{II}(\pi_1) =& \frac{1}{2}\pi_1\Gamma_P^{II}(1) + \frac{1}{2}\Gamma_P^{II}(\pi_1) + \frac{1}{2}(1-\pi_1)\Gamma_P^{II}(0)\\ =& -\frac{(1-\pi_1)(4-\pi_1)}{8(2-\pi_1)}.\end{aligned}$$

If type (iii) delegation equilibrium or the unique pure strategy communication equilibrium is played in the first period, the principal's expected payoff of the second period is

$$\begin{aligned}\hat{\Gamma}_P^{II}(\pi_1) =& \frac{1}{2}\pi_1\Gamma_P^{II}(1) + [\frac{1}{2}\pi_1 + (1-\pi_1)]\Gamma_P^{II}(\frac{\pi_1}{2-\pi_1})\\ =& -\frac{(2-\pi_1)(1-\pi_1)}{2(4-3\pi_1)}.\end{aligned}$$

Next, I investigate each type of equilibrium that can arise if the principal chooses delegation.

(i) $\sigma_G(0) < \sigma_G(1) = \sigma_B(0) < \sigma_B(1) = 1$. If this strategy profile constitutes an equilibrium, the following incentive constraints need to hold simultaneously:

$$-(\sigma_G(0)-0)^2+\Gamma_G(\pi_2(\sigma_G(0))) \geq -(0-0)^2+\Gamma_G(0) \tag{6}$$
$$-(\sigma_G(1)-1)^2+\Gamma_G(\pi_2(\sigma_G(1))) \geq -(1-1)^2+\Gamma_G(0) \tag{7}$$
$$-(\sigma_G(1)-1)^2+\Gamma_G(\pi_2(\sigma_G(1))) \geq -(\sigma_G(0)-1)^2+\Gamma_G(\pi_2(\sigma_G(0))) \tag{8}$$
$$\sigma_B(0)+\Gamma_B(\pi_2(\sigma_B(0))) \geq 1+\Gamma_B(0) \tag{9}$$
$$\sigma_B(0)+\Gamma_B(\pi_2(\sigma_B(0))) \geq \sigma_G(0)+\Gamma_B(\pi_2(\sigma_G(0))) \tag{10}$$
$$\sigma_B(1)+\Gamma_B(\pi_2(\sigma_B(1))) \geq \sigma_B(0)+\Gamma_B(\pi_2(\sigma_B(0))) \tag{11}$$
$$\sigma_B(1)+\Gamma_B(\pi_2(\sigma_B(1))) \geq \sigma_G(0)+\Gamma_B(\pi_2(\sigma_G(0))) \tag{12}$$

Inequality (6) requires the good agent has no incentive to deviate to action 0 which is his bliss point at state 0 and induces the principal to believe the agent is surely bad, i.e., $\pi_2(0) = 0$ (we assume the most pessimistic belief off the equilibrium path). If $\sigma_G(0)$ happens to be 0, the inequality still holds since $\Gamma_G(\pi_2(\sigma_G(0))) = \Gamma_G(1) > \Gamma_G(0)$. This constraint guarantees that any action between 0 and $\sigma_G(0)$ cannot be profitable for the good agent since they yield a lower current utility than action 0 but the same reputation as action 0. Clearly, it is not profitable for the agent to deviate to any action above $\sigma_G(0)$ at state 0 since it yields a lower current and also a lower reputation than action $\sigma_G(0)$. Similarly, inequality (7) and (8) require the good agent has no incentive to deviate to action 1 and $\sigma_G(0)$ at state 1, which ensures that it is not profitable to deviate to any other actions since they will yield either a lower current utility or a lower reputation, or both than action 1 or $\sigma_G(0)$. Analogous illustration for the bad agent's incentive constraints (9)-(12).

Recall that $\Gamma_G(\pi_2) = -\frac{1}{2}(\frac{1-\pi_2}{2-\pi_2})^2$ and $\Gamma_B(\pi_2) = \frac{1}{2-\pi_2}$. Moreover, in equilibrium we have $\pi_2(\sigma_G(0)) = 1$, $\pi_2(\sigma_G(1)) = \pi_2(\sigma_B(0)) = \pi_1$, and $\pi_2(\sigma_B(1)) = 0$. Plugging in and rearranging, we have

$$\text{(6) and (12)} \Rightarrow 0 \leq \sigma_G(0) \leq \frac{1}{2\sqrt{2}}, \tag{13}$$
$$\text{(9) and (11)} \Rightarrow \sigma_G(1) = \sigma_B(0) = \frac{3}{2} - \frac{1}{2-\pi_1}. \tag{14}$$

Given (13), we have (7) implies (8) so that we can ignore (8) since $-1 \leq -(\sigma_G(0) - 1)^2 \leq -(\frac{1}{2\sqrt{2}} - 1)^2 < -\frac{1}{8}$. Given (14), constraint (10) is equivalent to (12). Plugging (14) into (7), we have

$$0 < \pi_1 \leq \frac{4}{5}.$$

Therefore, if the prior π_1 is in $(0, 4/5]$, any strategy profile $(\sigma_G(0), \sigma_G(1), \sigma_B(0), \sigma_B(1))$ satisfying (13) and (14) constitutes an equilibrium in which the principal's expected payoff is

$$\begin{aligned}
& -\frac{1}{2}\pi_1(\sigma_G(0) - 0)^2 - \frac{1}{2}\pi_1(\sigma_G(1) - 1)^2 - \frac{1}{2}(1 - \pi_1)(\sigma_B(0) - 0)^2 \\
& \quad - \frac{1}{2}(1 - \pi_1)(\sigma_B(1) - 1)^2 \\
= & -\frac{1}{2}\pi_1\sigma_G(0)^2 + \frac{8\pi_1^3 - 33\pi_1^2 + 40\pi_1 - 16}{8(\pi_1 - 2)^2}.
\end{aligned}$$

Clearly, in the principal-optimal equilibrium, $\sigma_G(0)$ must be 0. We denote this maximum by

$$\Gamma_P^{(i)} = \frac{8\pi_1^3 - 33\pi_1^2 + 40\pi_1 - 16}{8(\pi_1 - 2)^2}, \ 0 < \pi_1 \leq \frac{4}{5}.$$

(ii) $\sigma_G(0) < \sigma_G(1) = \sigma_B(1) < \sigma_B(0) = 1$. If this strategy profile constitutes an equilibrium, the following incentive constraints need to hold simultaneously:

$$\begin{aligned}
& -(\sigma_G(0) - 0)^2 + \Gamma_G(\pi_2(\sigma_G(0))) \geq -(0 - 0)^2 + \Gamma_G(0) \\
& -(\sigma_G(1) - 1)^2 + \Gamma_G(\pi_2(\sigma_G(1))) \geq -(1 - 1)^2 + \Gamma_G(0) \\
& -(\sigma_G(1) - 1)^2 + \Gamma_G(\pi_2(\sigma_G(1))) \geq -(\sigma_G(0) - 1)^2 + \Gamma_G(\pi_2(\sigma_G(0))) \\
& \sigma_B(0) + \Gamma_B(\pi_2(\sigma_B(0))) \geq \sigma_B(1) + \Gamma_B(\pi_2(\sigma_B(1))) \\
& \sigma_B(0) + \Gamma_B(\pi_2(\sigma_B(0))) \geq \sigma_G(0) + \Gamma_B(\pi_2(\sigma_G(0))) \\
& \sigma_B(1) + \Gamma_B(\pi_2(\sigma_B(1))) \geq \sigma_B(0) + \Gamma_B(\pi_2(\sigma_B(0))) \\
& \sigma_B(1) + \Gamma_B(\pi_2(\sigma_B(1))) \geq \sigma_G(0) + \Gamma_B(\pi_2(\sigma_G(0)))
\end{aligned}$$

Similarly, we have if the prior π_1 is in $(0, 4/5]$, any strategy profile $(\sigma_G(0), \sigma_G(1), \sigma_B(0),$

$\sigma_B(1))$ satisfying

$$0 \le \sigma_G(0) \le \frac{1}{2\sqrt{2}}$$
$$\sigma_G(1) = \sigma_B(1) = \frac{3}{2} - \frac{1}{2-\pi_1}$$

constitutes an equilibrium in which the principal's expected payoff is

$$\begin{aligned} &-\frac{1}{2}\pi_1(\sigma_G(0)-0)^2 - \frac{1}{2}\pi_1(\sigma_G(1)-1)^2 - \frac{1}{2}(1-\pi_1)(\sigma_B(0)-0)^2 \\ &\quad -\frac{1}{2}(1-\pi_1)(\sigma_B(1)-1)^2 \\ =&-\frac{1}{2}\pi_1\sigma_G(0)^2 + \frac{4\pi_1^3 - 21\pi_1^2 + 32\pi_1 - 16}{8(\pi_1-2)^2}. \end{aligned}$$

Clearly, in the principal-optimal equilibrium, $\sigma_G(0)$ must be 0. We denote this maximum by

$$\Gamma_P^{(ii)} = \frac{4\pi_1^3 - 21\pi_1^2 + 32\pi_1 - 16}{8(\pi_1-2)^2},\ 0 < \pi_1 \le \frac{4}{5}.$$

By simple algebra, we find that $\Gamma_P^{(i)} > \Gamma_P^{(ii)}$ for all $\pi_1 \in (0, 4/5]$.

(iii) $\sigma_G(0) < \sigma_G(1) = \sigma_B(0) = \sigma_B(1)$. If this strategy profile constitutes an equilibrium, the following incentive constraints need to hold simultaneously:

$$-(\sigma_G(0)-0)^2 + \Gamma_G(\pi_2(\sigma_G(0))) \ge -(0-0)^2 + \Gamma_G(0) \tag{15}$$
$$-(\sigma_G(1)-1)^2 + \Gamma_G(\pi_2(\sigma_G(1))) \ge -(1-1)^2 + \Gamma_G(0) \tag{16}$$
$$-(\sigma_G(1)-1)^2 + \Gamma_G(\pi_2(\sigma_G(1))) \ge -(\sigma_G(0)-1)^2 + \Gamma_G(\pi_2(\sigma_G(0))) \tag{17}$$
$$\sigma_B(0) + \Gamma_B(\pi_2(\sigma_B(0))) \ge 1 + \Gamma_B(0) \tag{18}$$
$$\sigma_B(0) + \Gamma_B(\pi_2(\sigma_B(0))) \ge \sigma_G(0) + \Gamma_B(\pi_2(\sigma_G(0))) \tag{19}$$

Recall that $\Gamma_G(\pi_2) = -\frac{1}{2}(\frac{1-\pi_2}{2-\pi_2})^2$ and $\Gamma_B(\pi_2) = \frac{1}{2-\pi_2}$. Moreover, in equilibrium we have $\pi_2(\sigma_G(0)) = 1$, $\pi_2(\sigma_G(1)) = \pi_2(\sigma_B(0)) = \pi_2(\sigma_B(1)) = \frac{\pi_1}{2-\pi_1}$. Plugging in and rearranging, we have

$$(15) \Rightarrow 0 \le \sigma_G(0) \le \frac{1}{2\sqrt{2}}. \tag{20}$$

Given (20), we have (16) implies (17) so that we can ignore (17) since $-1 \leq -(\sigma_G(0)-1)^2 \leq -(\frac{1}{2\sqrt{2}}-1)^2 < -\frac{1}{8}$, and also (18) implies (19) so that we can ignore (19) since $\sigma_G(0)+1 < 3/2$. Moreover, we can reduce (16) and (18) to

$$(\sigma_B(0)-1)^2 + 2(\frac{1-\pi_1}{4-3\pi_1})^2 \leq \frac{1}{8} \tag{21}$$

$$\sigma_B(0) + \frac{2-\pi_1}{4-3\pi_1} \geq \frac{3}{2} \tag{22}$$

Therefore, any strategy profile $(\sigma_G(0), \sigma_G(1), \sigma_B(0), \sigma_B(1))$ satisfying (20)-(22) constitutes an equilibrium in which the principal's expected payoff is

$$\begin{aligned}
&-\frac{1}{2}\pi_1(\sigma_G(0)-0)^2 - \frac{1}{2}\pi_1(\sigma_G(1)-1)^2 - \frac{1}{2}(1-\pi_1)(\sigma_B(0)-0)^2 \\
&\quad - \frac{1}{2}(1-\pi_1)(\sigma_B(1)-1)^2 \\
=&-\frac{1}{2}\pi_1\sigma_G(0)^2 - \frac{1}{2}(\sigma_B(0)-1)^2 - \frac{1}{2}(1-\pi_1)\sigma_B(0)^2.
\end{aligned}$$

Clearly, in the principal-optimal equilibrium, $\sigma_G(0)$ must be 0. Define

$$\tilde{\Gamma}_P^{(iii)} = -\frac{1}{2}(\sigma_B(0)-1)^2 - \frac{1}{2}(1-\pi_1)\sigma_B(0)^2$$

where $\sigma_B(0)$ satisfies (21) and (22).

Note that the unconstrained maximizer of $\tilde{\Gamma}_P^{(iii)}$ is $\sigma_B(0) = \frac{1}{2-\pi_1} = a_1^P(\bar{m})$. It is straightforward to verify that if $2(4-\sqrt{2})/7 \leq \pi_1 < 1$, $\sigma_B(0) = \frac{1}{2-\pi_1}$ satisfies (21) and (22). Thus, $\sigma_B(0) = \frac{1}{2-\pi_1}$ is also the constrained maximizer. The corresponding maximum of the principal's expected payoff at this maximizer denoted by $\Gamma_P^{(iii)}$ is

$$\Gamma_P^{(iii)} = -\frac{\pi_1 - 1}{2(\pi_1-2)}.$$

Therefore, the delegation game can replicate the equilibrium outcome of the communication game when $2(4-\sqrt{2})/7 \leq \pi_1 < 1$. In particular, note that $\sigma_G(0) = a_1^P(\underline{m}) = 0$ and $\sigma_B(0) = a_1^P(\bar{m}) = \frac{1}{2-\pi_1}$ in equilibrium. Also recall that the principal obtains the same expected payoff of the second period, $\hat{\Gamma}_P^{II}$, from type (iii) equilibrium and the communication equilibrium.

On the other hand, if $\pi_1 < \frac{2(4-\sqrt{2})}{7}$, the unconstrained maximizer $\sigma_B(0) = \frac{1}{2-\pi_1}$ cannot be an equilibrium strategy since it does not satisfy constraint (22) any more.

Hence, the constrained maximum of the principal's expected payoff would be weakly lower in an type (iii) equilibrium of the delegation game than communication game. However, we need to check whether type (i) equilibrium can do better than communication or not.

By standard computation, we find that $\Gamma_P^{(i)} + \tilde{\Gamma}_P^{II} > \Gamma_P^{I} + \hat{\Gamma}_P^{II}$ if $\pi_1' < \pi_1 \leq 4/5$ and $\Gamma_P^{(i)} + \tilde{\Gamma}_P^{II} \leq \Gamma_P^{I} + \hat{\Gamma}_P^{II}$ if $0 < \pi_1 \leq \pi_1'$, where π_1' is determined by the equality $\Gamma_P^{(i)} + \tilde{\Gamma}_P^{II} = \Gamma_P^{I} + \hat{\Gamma}_P^{II}$. Moreover, it is straightforward to show $\pi_1' < \frac{2(4-\sqrt{2})}{7} < \frac{4}{5}$.

Combining the results above together, we can conclude in terms of the principal's expected payoff, if $\pi_1 \leq \pi_1'$, all equilibria of the delegation game are (weakly) worse than the equilibrium of the communication game. If $\pi_1' < \pi_1 \leq 4/5$, the principal obtains a higher expected payoff from the type (i) delegation equilibrium and therefore delegation is strictly better than communication. If $4/5 < \pi_1 < 1$, the type (iii) delegation equilibrium can replicate the communication equilibrium and therefore delegation is as good as communication.

By defining $\bar{\pi}_1$ in the proposition as π_1', we complete the proof.

□

This proposition shows that the optimal authority allocation in the first period depends on the principal's prior belief about the agent being good. Intuitively, when the fraction of the good agent is relatively high, it is less costly for the good agent to signal his type for a higher reputation in future. Also delegation makes the agent with authority behave more sensitively to the state when making decisions. And thus when the principal believes the agent is more likely to be aligned, she would prefer delegation over communication. However, when the fraction of the good agent is relatively low, the good agent needs to behave more aggressively to separate from the bad agent. This reputation effect makes the principal worse off from giving up control so that she would rather prefer to ask the agent for information.

3.5 Conclusion

In many situations, the motives of the expert may not be transparent to the decision maker. Also, the interactions between the decision maker and the expert are often repeated. In a long-run relationship, the expert may have an incentive to "look good," i.e., have reputational concerns. In such situations, if the decision maker can choose the decision making protocol, should she keep control and solicit information from the expert or delegate the decision making to the expert? I address this question in this paper.

I consider a two-period repeated game. In each period, the uninformed principal first decides whether to delegate the decision making to the informed agent who is either good (not biased) or bad (biased). If she does, the agent takes an action himself. If she does not, the agent sends a cheap talk message to the principal who then takes an action. I find that in the second period, the principal is better off by keeping control instead of delegating to the agent. In the first period, the communication equilibrium features the same structures as the second period. If the principal delegates control, the action that the agent takes not only affects his current utility but also signals his type, which affects his future utility. The optimal authority allocation depends on a prior cut-off. If the prior about the agent being good is above this cut-off, the principal prefers delegation over communication. Otherwise, communication dominates delegation from the principal's point of view.

Bibliography

Agastya, M., P. K. Bag, and I. Chakraborty (2014). Communication and authority with a partially informed expert. *The RAND Journal of Economics 45*(1), 176–197.

Aghion, P. and J. Tirole (1997). Formal and real authority in organizations. *Journal of Political Economy 105*(1), 1–29.

Aliprantis, C. D. and K. C. Border (2006). *Infinite Dimensional Analysis: A Hitchhiker's Guide* (3 ed.). Springer.

Alonso, R. and N. Matouschek (2007). Relational delegation. *The RAND Journal of Economics 38*(4), 1070–1089.

Alonso, R. and N. Matouschek (2008). Optimal delegation. *The Review of Economic Studies 75*(1), 259–293.

Amador, M. and K. Bagwell (2013). The theory of optimal delegation with an application to tariff caps. *Econometrica 81*(4), 1541–1599.

Ambrus, A. and S. Takahashi (2008). Multi-sender cheap talk with restricted state spaces. *Theoretical Economics 3*, 1–27.

Argenziano, R., S. Severinov, and F. Squintani (2016). Strategic information acquisition and transmission. *American Economic Journal: Microeconomics 8*(3), 119–155.

Armstrong, M. and J. Vickers (2010). A model of delegated project choice. *Econometrica 78*(1), 213–244.

Avery, C. and M. Meyer (2012). Reputational incentives for biased evaluators. *Unpublished manuscript*.

Battaglini, M. (2002). Multiple referrals and multidimensional cheap talk. *Econometrica 70*(4), 1379–1401.

Ben-Porath, E., E. Dekel, and B. L. Lipman (2014). Optimal allocation with costly verification. *The American Economic Review 104*(12), 3779–3813.

Benabou, R. and G. Laroque (1992). Using privileged information to manipulate markets: Insiders, gurus, and credibility. *The Quarterly Journal of Economics 107*(3), 921–958.

Bird, D. and A. Frug (2017). Dynamic nonmonetary incentives. *Unpublished manuscript*.

Bubb, R. and P. L. Warren (2014). Optimal agency bias and regulatory review. *Journal of Legal Studies 43*, 95–135.

Carroll, G. and G. Egorov (2017). Strategic communication with minimal verification. Working paper.

Chakraborty, A. and R. Harbaugh (2007). Comparative cheap talk. *Journal of Economic Theory 132*(1), 70–94.

Chakraborty, A. and R. Harbaugh (2010). Persuasion by cheap talk. *The American Economic Review 100*(5), 2361–2382.

Che, Y.-K. and N. Kartik (2009). Opinions as incentives. *Journal of Political Economy 117*(5), 815–860.

Crawford, V. P. and J. Sobel (1982). Strategic information transmission. *Econometrica: Journal of the Econometric Society*, 1431–1451.

Deimen, I. and D. o. Szalay (2019). Delegated expertise, authority, and communication. *American Economic Review 109*(4), 1349–1374.

Demski, J. S. and D. E. M. Sappington (1987). Delegated expertise. *Journal of Accounting Research 25*(1), 68–89.

Dessein, W. (2002). Authority and communication in organizations. *The Review of Economic Studies 69*(4), 811–838.

Ely, J. C. and J. Välimäki (2003). Bad reputation. *The Quarterly Journal of Economics 118*(3), 785–814.

Erlanson, A. and A. Kleiner (2017). Costly verification in collective decisions. Working paper.

Garfagnini, U., M. Ottaviani, and P. N. Sørensen (2014). Accept or reject? an organizational perspective. *International Journal of Industrial Organization 34*, 66–74.

Gentzkow, M. and J. M. Shapiro (2006). Media bias and reputation. *Journal of political Economy 114*(2), 280–316.

Gerardi, D. and L. Yariv (2008). Costly expertise. *American Economic Review: Papers & Proceedings 98*(2), 187–193.

Glazer, J. and A. Rubinstein (2004). On optimal rules of persuasion. *Econometrica 72*(6), 1715–1736.

Golosov, M., V. Skreta, A. Tsyvinski, and A. Wilson (2014). Dynamic strategic information transmission. *Journal of Economic Theory 151*, 304–341.

Grossman, S. J. (1981). The informational role of warranties and private disclosure about product quality. *The Journal of Law and Economics 24*(3), 461–483.

Grossman, S. J. and O. D. Hart (1986). The costs and benefits of ownership: A theory of vertical and lateral integration. *Journal of Political Economy 94*(4), 691–719.

Guo, Y. (2016). Dynamic delegation of experimentation. *American Economic Review 106*(8), 1969–2008.

Guo, Y. and J. Hörner (2015). Dynamic mechanisms without money. Technical report, Reihe Ökonomie/Economics Series, Institut für Höhere Studien (IHS).

Hart, O. and J. Moore (1990). Property rights and the nature of the firm. *Journal of Political Economy 98*(6), 1119–1158.

Hart, S., I. Kremer, and M. Perry (2017). Evidence games: truth and commitment. *The American Economic Review 107*(3), 690–713.

Holmström, B. (1977). *On Incentives and Control in Organizations*. Ph.D. dissertation, Stanford University.

Holmström, B. (1984). On the theory of delegation. In M. Boyer and R. Kihlstrom (Eds.), *Bayesian Models in Economic Theory*, pp. 115–141. North-Holland.

Kamenica, E. and M. Gentzkow (2011). Bayesian persuasion. *The American Economic Review 101*(6), 2590–2615.

Kartik, N. (2009). Strategic communication with lying costs. *The Review of Economic Studies 76*(4), 1359–1395.

Levy, G. (2004). Anti-herding and strategic consultation. *European Economic Review 48*(3), 503–525.

Levy, G. and R. Razin (2007). On the limits of communication in multidimensional cheap talk: a comment. *Econometrica 75*(3), 885–893.

Li, H. (2001). A theory of conservatism. *Journal of Political Economy 109*(3), 617–636.

Li, J., N. Matouschek, and M. Powell (2017). Power dynamics in organizations. *American Economic Journal: Microeconomics 9*(1), 217–41.

Lipnowski, E. and J. Ramos (2018). Repeated delegation. *Unpublished manuscript*.

Lipnowski, E. and D. Ravid (2017). Cheap talk with transparent motives. Working paper.

Malcomson, J. M. (2009). Principal and expert agent. *B.E. Journal of Theoretical Ecoonomics 9*(1), Article 17.

Martimort, D. and A. Semenov (2006). Continuity in mechanism design without transfers. *Economics Letters 93*(2), 182–189.

Melumad, N. D. and T. Shibano (1991). Communication in settings with no transfers. *The RAND Journal of Economics*, 173–198.

Milgrom, P. R. (1981). Good news and bad news: Representation theorems and applications. *The Bell Journal of Economics*, 380–391.

Morgan, J. and P. C. Stocken (2003). An analysis of stock recommendations. *RAND Journal of economics*, 183–203.

Morris, S. (2001). Political correctness. *Journal of political Economy 109*(2), 231–265.

Mylovanov, T. and A. Zapechelnyuk (2017). Optimal allocation with ex post verification and limited penalties. *The American Economic Review 107*(9), 2666–94.

Ottaviani, M. (2000). The economics of advice. *University College London, mimeo*.

Ottaviani, M. and P. N. Sørensen (2006a). Professional advice. *Journal of Economic Theory 126*(1), 120–142.

Ottaviani, M. and P. N. Sørensen (2006b). Reputational cheap talk. *The Rand journal of economics 37*(1), 155–175.

Prat, A. (2005). The wrong kind of transparency. *American economic review 95*(3), 862–877.

Prendergast, C. and L. Stole (1996). Impetuous youngsters and jaded old-timers: Acquiring a reputation for learning. *Journal of political Economy 104*(6), 1105–1134.

Rush, A., V. Smirnov, and A. Wait (2010). Communication breakdown: Consultation or delegation from an expert with uncertain bias. *The BE Journal of Theoretical Economics 10*(1).

Scharfstein, D. S. and J. C. Stein (1990). Herd behavior and investment. *The American Economic Review*, 465–479.

Semenov, A. (2018). Delegation to a possibly ignorant agent. *Canadian Journal of Economics/Revue canadienne d'économique 51*(1), 64–93.

Sobel, J. (1985). A theory of credibility. *The Review of Economic Studies 52*(4), 557–573.

Spence, M. (1973). Job market signalling. *Quarterly Journal of Economics 87*(3), 355–374.

Szalay, D. (2005). The economics of clear advice and extreme options. *The Review of Economic Studies 72*(4), 1173–1198.

Xu, H. (2012). Reputational concern with endogenous information acquisition. *Unpublished manuscript*.

Curriculum Vitae

Ph.D.	Economics	University of Mannheim, 2019
M.A.	Economics	Fudan University, 2013
B.A.	Economics	Zhongnan University of Economics and Law, 2010